

Retrofitting for the Logistics Revolution

*Designing Marine Corps
Inventories to Support
the Warfighter*

RAND

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Ronald D. Fricker, Jr.

Marc L. Robbins

DISC INSPECTED
20000720 098

The research described in this report was sponsored by the United States Marine Corps (USMC). The research was conducted in RAND's National Defense Research Institute, a federally funded research and development center supported by the Office of the Secretary of Defense (OSD), the Joint Staff, the unified commands, and the defense agencies under Contract DASW01-95-C-0059.

Library of Congress Cataloging-in-Publication Data

Fricker, Ronald D., 1960–

Retrofitting for the logistics revolution : designing Marine Corps inventories to support the warfighter / Ronald D. Fricker, Jr., Marc Robbins.

p. cm

“MR-1096-USMC.”

Includes bibliographical references.

ISBN 0-8330-2855-3

1. United States. Marine Corps—Supplies and stores. 2. United States. Marine Corps—Management—Data processing. 3. United States. Marine Corps—Management—Automation. I. Robbins, Marc L., 1954– II. Title.

VE480 .F75 2000

359.9'6621'0973—dc21

00-034149

RAND is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND® is a registered trademark. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

© Copyright 2000 RAND

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

Published 2000 by RAND

1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1200 South Hayes Street, Arlington, VA 22202-5050

RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information,
contact Distribution Services: Telephone: (310) 451-7002;

Fax: (310) 451-6915; Internet: order@rand.org

PREFACE

In 1996 Marine Corps leaders, recognizing the need to improve logistics processes and supply support to the Fleet Marine Force (FMF), implemented the Precision Logistics program. Precision Logistics draws on much of the process improvement philosophy and many of the methodologies employed by commercial industry to achieve significant, and often radical, performance improvement over the past two decades.

Initial Precision Logistics efforts focused on defining and measuring logistics processes, the first two steps in a “define-measure-improve” continuous improvement methodology. Major Precision Logistics initiatives included the measurement and dissemination of order and ship times (OSTs) and repair cycle times (RCTs). Recently, the Marine Corps has also introduced a metric directly related to measuring repair parts inventory performance, the “ERO fill rate,” to promote improvements in retail repair parts inventory. In general, Precision Logistics seeks improvement in the “nuts and bolts” of support to Marine Corps customers: requisition submission and processing, distribution, repair processes, inventory management, and so forth.

More recently, the Marine Corps has initiated the Integrated Logistics Capability (ILC) Initiative to plan the next generation of core logistics processes and the supporting information technology infrastructure (USMC, 1999). The stated goals are to evaluate and adapt best business practices, redefining Marine Corps logistics processes as a result, and then specifying a plan and actions to achieve the desired future state.

This report presents new inventory management tools and techniques, relevant to both Precision Logistics and the ILC Initiative, to help the Marine Corps (and other organizations involved in inventory management) continue to improve logistics and repair processes. We advocate specific inventory management methodology improvements and demonstrate that these improvements lead directly to improved inventory performance.

Our methods for improving how the Marine Corps manages inventory are based on modern and novel computational methods that are just now possible with today's powerful desktop computers. We take the position that these improvements will dramatically improve supply system performance and demonstrate the level of expected improvement for the Marine Corps using actual data. Supply system performance improvements include reduced OST and RCT, and increased ERO fill rates, as well as improvements in more traditional inventory metrics, such as fill rate. While our recommended changes are projected to result in a more fiscally efficient operation, our main focus is on improving retail supply from the central manager of inventory in the Marine Corps to the ultimate user—the *warfighter*.

Our work was sponsored by the Deputy Chief of Staff for Installations and Logistics (DC/S I&L). It should be of interest to Marine Corps logisticians, supply and parts managers, and others interested in logistics, as well as those supporting repair processes, and quality practitioners involved with performance measurement and process improvement in general. While the results in this report are couched in Marine Corps-specific terminology, all of our methods are founded in standard inventory theory and are directly applicable to any facility that maintains an inventory. Comments about the report may be sent to the authors by e-mail (Ron_Fricker@rand.org or Marc_Robbins@rand.org), or by post at RAND, 1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138.

The research was conducted within the Forces and Resources Policy Center of RAND's National Defense Research Institute, a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, and the defense agencies.

CONTENTS

Preface	iii
Figures	vii
Tables	ix
Summary	xi
Acknowledgments	xix
Acronyms	xxi
Chapter One	
BUILDING A MECHANISM FOR CHANGE: PRECISION LOGISTICS AND THE INTEGRATED LOGISTICS CAPABILITY INITIATIVE	1
Outline of the Report	5
Chapter Two	
THE INVENTORY FULCRUM: MARINE CORPS LOGISTICS POLICY AND PRACTICE	7
Repair and Supply in the Fleet Marine Force	7
Supply	8
Inventory Performance Metrics	10
Looking into the Mirror of Current Performance	11
RCT Performance in the FMF	15
The ERO Fill Rate: Definition and Performance	18
Chapter Three	
A LEVER FOR RAISING PERFORMANCE: AN IMPROVED INVENTORY METHODOLOGY	21

Illustrating Problems with Current Marine Corps	
Inventory Methodology	23
Properly Calculating the ROP: The Bootstrap ROP	24
How the Bootstrap ROP Methodology Works	25
Application to USMC Data	28
Examples of Bootstrap ROPs	30
Incorporating "Bang for the Buck": Dollar Banding	33
A Simple Example	33
Combining the Dollar-Banding and Bootstrap ROP	
Methodologies	36
Reducing the Length and Variability of the Replenishment	
Process	37
Conclusions	42
Chapter Four	
PUSHING THE LEVER: DEMONSTRATIONS OF	
IMPROVED INVENTORY PERFORMANCE	45
The Virtual SMU	45
Creating a New GABF Using the Bootstrap	46
Further Improvements with Dollar Banding	49
Conclusions	50
Chapter Five	
FUTURE STEPS: IMPLEMENTATION,	
METHODOLOGICAL IMPROVEMENTS,	
GENERALIZATIONS	53
For the Marine Corps	54
For Other Organizations	56
Appendix	
A. CRITICAL REPAIRS AND THE ERO FILL RATE	59
B. BASIC INVENTORY THEORY AND TERMINOLOGY	65
C. COMMON MARINE CORPS REORDER POINT AND THE	
OPERATING LEVEL MISCONCEPTIONS	71
D. USING THE BOOTSTRAP TO CALCULATE REORDER	
POINTS	75
E. THE VIRTUAL SMU	87
References	99

FIGURES

S.1 Logistics Process Improvements Create Potential for Synergies	xii
1.1 The “Define-Measure-Improve” Methodology	3
1.2 Logistics Process Improvements Create Potential for Synergies	4
2.1 Impact on RCT of Parts Supplied from Wholesale	14
3.1 The Performance of Hypothetical Item 1 Using the Current Marine Corps Methodology for Setting the RO and ROPs	24
3.2 The Performance of Hypothetical Item 2 Using the Current Marine Corps Methodology for Setting the RO and ROPs	25
3.3 Distribution of Replenishment OSTs for I MEF SMU (MMC100), CY 1998	38
3.4 Distribution of Replenishment OSTs for II MEF SMU (MML100), CY 1998	39
3.5 Distribution of Replenishment OSTs for III MEF SMU (MMR100), CY 1998 (Change in Median and 75th Percentiles Between First and Second Half of Year Noted)	41
B.1 An Illustration of the Relationship of the IP to the On-Hand Stock, the RO, and the ROP	67
B.2 An Illustration of the Effect of Increasing the ROP from “S” in Figure B.1 to “S+d”	68
B.3 Lead Time Indicated by the Gray Arrows	69
E.1 Schematic Outline of the Virtual SMU	90
E.2 Dramatically Improved Fill Rate Performance Under a Simple Buy Policy	96

E.3 Performance of the Inventory Under the Simple Buy Policy	97
---	----

TABLES

S.1	Comparison of the Projected Inventory Performance Using Actual and New Techniques	xvi
2.1	Number of Requisitions per ERO, FMF September 1998	12
2.2	Illustrations of Last Part Received	13
2.3	Examples of Inexpensive Parts Driving RCT	16
2.4	Principal End Item Repair Cycle Time—CY 1998	17
2.5	Principal End Item Average Time in Repair Status—CY 1998	17
2.6	Secondary Repairable Repair Cycle Time—CY 1998	17
2.7	Secondary Repairable Average Time in Repair Status—CY 1998	18
2.8	ERO Fill Rate Performance by MEF (CY 1998): Overall, by PEIs, and by Secondary Repairables	19
3.1	Comparison of the Actual April 27 GABF to Various Alternative Bootstrap GABFs	29
3.2	Various NSNs and Their Annual Demand Pattern for One Year with Three Possible Bootstrap ROPs Corresponding to Risks of 0.01, 0.05, and 0.10	31
3.3	Comparison of the Various Dollar-Banding Schemes with the Demand-Only Baseline that Follows USMC Guidelines	36
3.4	Comparison of Bootstrap Inventories with Different Lead-Time Distributions	40
4.1	Comparison of the Actual GABF Performance with Various Bootstrap GABFs	48
4.2	Comparison of the Performance of Various Bootstrap GABFs	50

x Retooling for the Logistics Revolution

D.1 Comparison of the Probability of Not Stocking-Out (1-r) with the Actual Observed Rate for Five 30-Day Periods for All Lines (NSNs) that Had a Positive Bootstrap ROP	81
---	----

SUMMARY

In 1996 Marine Corps leaders, recognizing the need to improve logistics processes and supply support to the Fleet Marine Force (FMF), implemented the Precision Logistics program. One of the key ideas of Precision Logistics is envisioning logistics as a set of integrated processes: actions occur in a linked chain, with outputs from one source being transformed into inputs that satisfy the next link in the chain. Logistics then becomes an unbroken chain of activities, linking many actors, all aimed at achieving an overarching goal of providing efficient and effective warfighter support to sustain military operations in continuously faster, better, and cheaper fashion. While this concept is expressed in military-specific terminology, the idea is just as applicable to many civilian manufacturing and repair operations and is often referred to as "supply chain management" in the current logistics literature.

To date, the Marine Corps has emphasized improving performance in supply distribution processes (order and ship times), both for retail (local) and wholesale (nonlocal) stock, and repair cycle time. Marines now have tools in place to help show performance trends and diagnostic tools to help explain reasons for poor performance (Robbins et al., 1998). Another critical area that Marine units have begun attacking is inventory policy and execution. In addition, the Marine Corps has initiated the Integrated Logistics Capability (ILC) Initiative to plan the next generation of core logistics processes and the supporting information technology infrastructure. A more effective set of inventory policies and procedures, backed by metrics, performance reports, and diagnostic tools, could be an important

lever for obtaining significantly more effective logistics performance. Figure S.1 states this point schematically.

This figure conveys the idea that logistics processes are not isolated and that improvements in the performance of one can pave the way to superior performance in another (and alternatively, poor performance in one can hamper performance in related processes). Improved order and ship and repair processes can facilitate an improved inventory determination process, in both cases by reducing pipeline times and variability and thus the amount of safety stock needed to cover future demands. Alternatively, an improved inventory position could, with fewer resources invested, result in reduced order and ship time (OST) (because more items will be available locally) and decreased repair cycle time (because less time will be lost waiting for parts).

Baselining performance can motivate the need for change. Recent evidence of shortfalls in Marine Corps retail inventory makes such a case. While the *fill rate* (percentage of demands satisfied from

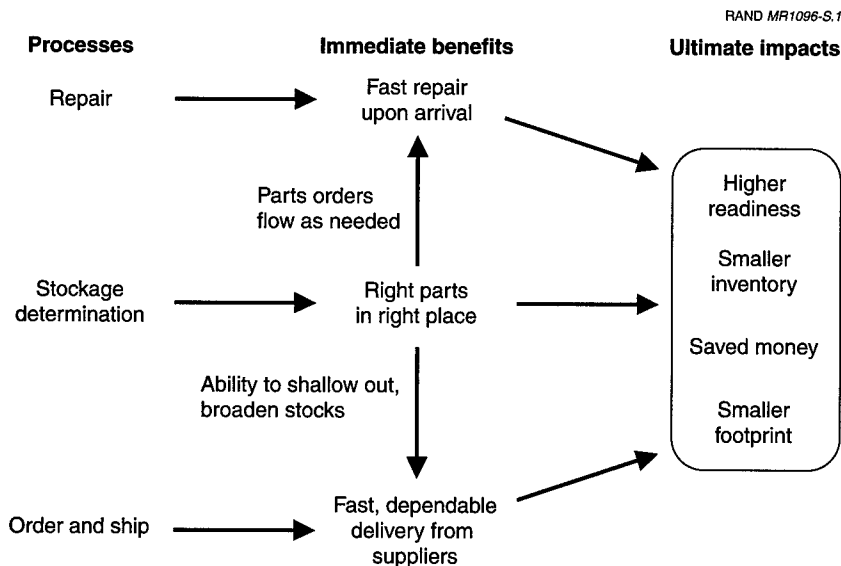


Figure S.1—Logistics Process Improvements Create Potential for Synergies

local inventory) would appear reasonably high (often in the 60–70 percent region), service to the customer, from the perspective of that customer, paints a different picture. More than half of necessary repair actions must go outside local Marine supply points, to the wholesale supply system run by the Defense Logistics Agency, for parts critical to those repairs. For certain types of repairs, such as for major subassemblies, that figure is in the 70–80 percent range. And often the part required is cheap. Lack of a \$1 part will hold down a deadlined weapon system even as other parts—worth thousands of dollars—are immediately available. These weaknesses in supply performance are a major contributor to poor repair times, which tend to be both lengthy on average and highly variable. Increased downtime for weapon systems reduces readiness, and delays in repairing major subassemblies lengthen the order and ship time for those items when required.

This report advocates specific supply and inventory management improvements to deal with these weaknesses. In particular, we present methods for improving how the Marine Corps stocks inventory—methods based on modern computational methods now possible with the advent of today’s powerful desktop computers. In keeping with the ILC Initiative, these changes are on the cutting edge of logistics practices. With them, we demonstrate significant inventory performance improvement in such conventional measures as fill rate but also in measures that go directly to what the supply system’s customers really care about, such as the ability of local supply to deliver *all* the parts needed for critical repairs (which we measure using a metric called the *ERO fill rate*). A better supply system will lead directly to reduced repair cycle time and thus to significantly improved support to the warfighter and lower system costs.

REVOLUTIONIZING INVENTORY MANAGEMENT IN THE MARINE CORPS

Local inventory in the Fleet Marine Force is held by a central manager, the SASSY Management Unit,¹ or SMU. Managing an SMU inventory is a three-part problem. First, the inventory manager must

¹“SASSY” stands for Supported Activities Supply System.

decide what to stock. This decision can be based on a variety of considerations, such as past demand history or the importance of the item. Second, having decided to stock an item, the inventory manager must set a reorder point (ROP). This is a decision based on an assessment of risk: What risk of stock-out is the inventory manager willing to assume when it is time to submit a replenishment order? The higher the ROP is set the lower the risk of running out during the replenishment time, but the more capital is tied up in safety stock. Third and finally, given an ROP, the inventory manager must set an operating level, which defines how often replenishment requisitions must be placed. Setting the operating level is a trade-off between the frequency of submitting replenishment orders and stock storage capacity.

We present two approaches to improve the Marine Corps retail inventory performance: “bootstrapping” the ROP, and “dollar banding” the inventory. These techniques concentrate on the first two aspects of the three-part problem, and implemented together, would revolutionize inventory management in the Marine Corps.

- *Bootstrapping* the ROP is a state-of-the-art statistical technique that allows the inventory manager to set ROPs by risk of stock-out. This approach automatically sizes the inventory to account for variations in demand—variation inherent in the supply system and not accounted for using the current methodology.
- *Dollar banding* incorporates both cost and demand in the stock decision. Essentially, it operates on the principle that many more inexpensive items can be stocked for the cost of very few expensive items. We will show that both the bootstrap ROP and dollar banding can be synergistically combined.

MEASURING THE IMPACT OF SUPPLY PERFORMANCE ON READINESS: THE ERO FILL RATE

The ERO fill rate is a measure of inventory effectiveness from the perspective of the mechanic, i.e., how many repairs receive *all* their high-priority parts from local supply. It has an intuitively direct connection to unit readiness. Yet, for all of its intuitive connection to readiness, Marine Corps supply personnel have little direct feedback about how supply policies connect to the ERO fill rate. Current

inventory practices were designed more for improving standard metrics, such as fill rate, whereas improvement in the ERO fill rate will require some significant departures from these traditional techniques. Although additional research is required to develop a methodology designed specifically for the ERO fill rate, immediate improvement is possible with the application of the bootstrap and dollar-banding techniques. As we will demonstrate, these techniques allow the Marine Corps to properly stock many more items for the same inventory investment, which results in dramatic ERO fill rate improvements. As Chapter Four will show, we apply these techniques, and others, to define new inventories that significantly outperform existing inventories. In particular, we proposed and evaluated through simulation new inventories that

- perform as well as the existing inventory in terms of fill rate, satisfaction rate, accommodation rate, and ERO fill rate—at less than one-half the current inventory cost, and,
- outperform the existing inventory, achieving ERO fill rates of more than 70 percent and fill rate metrics near or better than 90 percent at three-fourths of the current inventory value.

We note that a reduction in wholesale OST would facilitate the implementation of these new strategies, as well as improve current system performance. Our simulations show that the current long and variable OSTs experienced by the Marine Corps hamper improvement in inventory performance. Indeed, it is our position that *dramatic* inventory performance improvement is achievable only in conjunction with marked reductions in OSTs.

DEMONSTRATING THE BENEFIT TO THE CUSTOMER: THE WARFIGHTER

We both discuss and apply the techniques using a new analytical tool: the Virtual SMU. RAND developed the Virtual SMU to test various policy alternatives “off-line” but under the rigors of actual demand. Its purpose is to evaluate how new inventory management techniques will perform in practice, with the goal of selecting inventory management policies, techniques, and algorithms that perform well. We use it to answer the question, “How can the SMU achieve the best possible *overall* performance within imposed funding con-

straints?" The Virtual SMU works by "playing back" actual demand histories experienced by Marine units, so that all the complexities of the real world are experienced.

Table S.1 compares the performance of an actual inventory in a major Marine organization (the 1st Marine Expeditionary Force [MEF] at Camp Pendleton, California) using actual operating parameters with our proposed new inventory created using the bootstrap and dollar-banding techniques. Here we see that the new techniques are projected to achieve a fill rate of 87 percent and an ERO fill rate of 75 percent. These are, of course, simulation results, though they are based on actual data and come as close to the real world as possible (without actually implementing the techniques in an MEF). We expect that similar results can be achieved in practice.

Table S.1 highlights the potential large benefits as well as the change in management philosophy required. The standard measure, fill rate, is much higher, and so is the measure of true customer satisfaction, the ERO fill rate. The investment in inventory is much less, dropping some 29 percent while achieving significantly better performance. While the dollar value of stock held is much less, the variety of stock almost triples, with the number of lines (specific types of items) increasing from 13,159 to 32,537.

Better performance is also enhanced by less variable and significantly reduced OSTs. Reduced OST frees up capital previously required for safety stock (to cover the inventory when a random, long

Table S.1
Comparison of the Projected Inventory Performance Using Actual and New Techniques

	Projected Inventory Performance	
	Actual Marine Corps Guidelines	New Techniques
ERO fill rate	54%	75%
Fill rate	72%	87%
Inventory value at requisitioning objective	\$24 million	\$17 million
Number of lines stocked	13,159	32,537

OST occurs). This capital can then be applied to stocking many more items, particularly those that are cheap and have low demand (so they do not require frequent replenishment). In addition, the depth of stocks is sized to fit demand history and avoid back-orders, while permitting a great reduction in the value of the inventory required. The result is an inventory that more effectively supports the Marine Corps warfighter.

BETTER INVENTORY MANAGEMENT IS INCREASINGLY IMPORTANT

The changes we propose are relevant to both the Precision Logistics program and to the ILC Initiative. In particular, new metrics, such as ERO fill rate and repair cycle time (RCT), can guide process improvement under Precision Logistics and form the basis for a new suite of integrated metrics in the ILC Initiative. The new inventory management techniques can similarly be applied as incremental improvements under Precision Logistics as well as becoming the standard in a future inventory management system resulting from the ILC. While the ILC Initiative proposes to reduce the "iron mountain" using information technology, some part of that mountain will always remain. Thus, cutting-edge inventory techniques, such as the bootstrap ROP, will be required to effectively manage even an "iron molehill." Indeed, the smaller the mountain, the more critical it will be to manage it effectively.

ACKNOWLEDGMENTS

This report represents a joint effort between RAND researchers and a large number of Marine logisticians involved in performance measurement and committed to process improvement.

We wish to thank our project sponsors, Major General Geoffrey Higginbotham and Major General Joseph Stewart, both previous Deputy Chiefs of Staff for Installations and Logistics (DC/S I&L), for their strong efforts in supporting our work. We also thank Lieutenant General Gary McKissock for his years of support of RAND efforts to help improve Marine logistics, both as commanding general of the 1st Force Service Support Group (FSSG) and the Marine Corps Materiel Command and as Deputy Commandant for Installations and Logistics.

At DC/S I&L we are grateful for the help we have received from many people, including Colonel Ed Dillard, our project monitor, and Major George Pointon, Major Tom Keating, and Captain Chris Goodhart. Captain Goodhart's contributions deserve special recognition. Chris assisted us—indeed collaborated with us—in numerous ways, including aiding in the literature review, evaluating the technical merit of the methodology, and carefully proofreading and commenting on much of this work.

At 1st FSSG, Camp Pendleton, we received much expert advice and significant assistance from Brigadier General Rick Kelly and his staff, including Lieutenant Colonel Jim Kessler, Major Doug Turlip, Major Terry Flannery, and CWO2 Barry Newland. At II MEF, we appreciate the help and input we received from Lieutenant Colonel Roy Truba and Major Frank Payne.

We also benefited from the strong and continuing support of Colonel Jerry Varela and the entire staff of the U.S. Marine Corps Materiel Command. In particular, we are indebted to the Marine Data Analysis Center where we received invaluable data and support from Mike Carroll. We also wish to thank all the other Marines and Marine civilians who helped make this a better product.

This work was strengthened as a result of collaboration with our RAND colleagues working on the Army's Velocity Management program, in particular, John Dumond, Rick Eden, Ken Girardini, Mark Totten, and Sung-Ho Ahn. We would also like to thank Kevin Gue of the Naval Postgraduate School and Lou "Kip" Miller of RAND, who reviewed this report. They provided many helpful and insightful comments that substantially improved both the content and presentation of the report. This sometimes complicated document is much less so thanks to the efforts of Donna Keyser, who helped us streamline the presentation and make the story clearer, and Jerene Kelly, who made the text less assaulting on the eye. Finally, we would be terribly remiss if we failed to acknowledge the outstanding programming and database management support of Pat Boren. Without her skillful programming, insight, and dedication to detail none of this would have been possible. All errors and omissions are solely our responsibility.

ACRONYMS

AAV	Amphibious assault vehicle
ATLASS	Asset Tracking for Logistics and Supply System
BO	Back-order
CATCODE	Maintenance Category Code
CATMDAYS	Days in catcode M (deadlined-MCGERR) Status
CEC	Combat Essentiality Code
D_t	Number of Demands for day t
DC/S I&L	Deputy Chief of Staff (Installation and Logistics)
DI	Due-in (stock)
DLA	Defense Logistics Agency
DO	Due-out (stock)
DoD	Department of Defense
DOS	Days of supply
DRIS	Date Received in Shop
EOQ	Economic order quantity
ERO	Equipment Repair Order
EROSL	Equipment Repair Order Shopping List
FMF	Fleet Marine Force
FSC	Federal Supply Class

FSSG	Force Service Support Group
FY98	Fiscal Year 1998 (October 1997 to September 1998)
GABF	General Account Balance File
ICP	Inventory control point
IIP	Initial Issue Provisioning
ILC	Integrated Logistics Capability (Initiative)
I MEF	1st Marine Expeditionary Force
II MEF	2nd Marine Expeditionary Force
III MEF	3rd Marine Expeditionary Force
IP	Inventory position
IP _i	IP for NSN <i>i</i>
<i>L</i>	Lead Time
LTD	Lead-Time Demand
MARCOR- LOGBASE	Marine Corps Logistics Base
MARDIV	Marine Division
MCLB	Marine Corps Logistics Base
MCGERR	Marine Corps Ground Equipment Resource Reporting
MCO	Marine Corps Order
MEF	Marine Expeditionary Force
MEU	Marine Expeditionary Unit
MILSTRIP	Military Standard Requisitioning and Issue Procedures
MIMMS	Marine Corps Integrated Maintenance Management System
MRO	Materiel release order
NIIN	National Item Identification Number
NCO	Noncommissioned officer
NSN	National Stock Number

NORS	Not operational for reason of supply
OH	On-hand (stock)
OL	Operating Level
OST	Order and ship time
PEBs	Preexpended bins
PEI	Principal end item
PL	Precision Logistics
PLT	Procurement Lead Time
r	Risk (as a probability)
RCT	Repair cycle time
RIC	Routing Identifier Code
RIP	Repairable issue point
RO	Requisitioning objective
ROP	Reorder point
SAC	Storage Activity Code
SASSY	Supported Activities Supply System
S_L	Sum of Demands Occurring in L Successive Time Periods
SL	Safety Level
SMU	SASSY Management Unit
t	Time Subscript
TAM	Table of Authorized Materiel
TAMCN	Table of Authorized Materiel Control Number
USMC	U.S. Marine Corps
VM	Velocity Management

**BUILDING A MECHANISM FOR CHANGE: PRECISION
LOGISTICS AND THE INTEGRATED LOGISTICS
CAPABILITY INITIATIVE**

From the Department of Defense level down to that of each of the separate services, a new paradigm is coming into focus in how logistics needs to be executed to support the warfighter in the new deployment-oriented, reduced-budget environment. This new approach is couched in a seeming plethora of new names—"Integrated Logistics Capability," "Velocity Management," "Focused Logistics," "Agile Combat Support," and so forth. Behind each of these terms, and the structures and initiatives underlying them, is the idea that logistics systems once built for stable, forward-located theaters must be rebuilt for the new environment. To that end, the military has begun looking outside for good ideas to apply, especially from the emerging revolution in commercial logistics. One critical idea is that all logistics actions are connected, tying customer and supplier, and that feedback on the output of logistics—how efficiently and effectively support is delivered to the warfighter—must flow smoothly to all the actors in the process.

The Marines' answers to this challenge are called "Precision Logistics" and the "Integrated Logistics Capability (ILC) Initiative." The Marine Corps initiated the ILC Initiative to plan the next generation of core logistics processes and the supporting information technology infrastructure (USMC, 1999). The stated goals are to evaluate and adapt best business practices, redefining Marine Corps logistics processes as a result, and then to specify a plan and actions to achieve the desired future state. Precision Logistics refers less to how

logistics is done than to how one *thinks* about how logistics is done, and the consequent translation of that thought to action. It demands “a cultural and paradigm change in the way [Marines] think and operate.” (Krulak, 1997.) The core of this change is to think about logistics as a set of integrated processes, in which actions occur on a linked chain, with inputs from one source being transformed into outputs that satisfy the next link on that chain. Logistics then becomes part of a seamless web of activities, linking many actors, all aimed at achieving an overarching goal of providing effective and required support to the warfighter to sustain military operations in continuously faster, better, and cheaper fashion.

Precision Logistics rests, at its core, on a simple methodology—“Define-Measure-Improve” (see Figure 1.1).¹ A logistics process (repairing a broken vehicle, ordering and receiving gear) must be understood beginning to end, without gaps, and the role of all relevant actors must be comprehended by each of those actors. Process performance must be *measured* to the extent possible, because systemic knowledge and the drive to improvement cannot be based on often misleading anecdotal evidence. Full understanding of process structure and performance pave the way to *improvement*, where the process owners (the supply clerks, mechanics, their noncommissioned officers, and on up through the commanding officers) who best know the process and the fixes required are empowered by the command structure to try out new ideas, weed out bad ones, encourage good ones, and reiterate the drive to get constantly better logistics processes.

The Marine Corps, to date, has put major emphasis on improving performance in supply distribution processes (“order and ship” times), both for retail (local) and wholesale (nonlocal) stock, and repair (repair cycle time). Marines now have tools in place to help show performance trends and diagnostic tools to help explain poor

¹For additional discussion on the Define-Measure-Improve methodology, as well as information on its successful implementation in the Army’s Velocity Management program, see Dumond, Eden, and Folkesson (1995), Edwards and Eden (1998), and “Speeding” (1998).

RAND MR1096-1.1

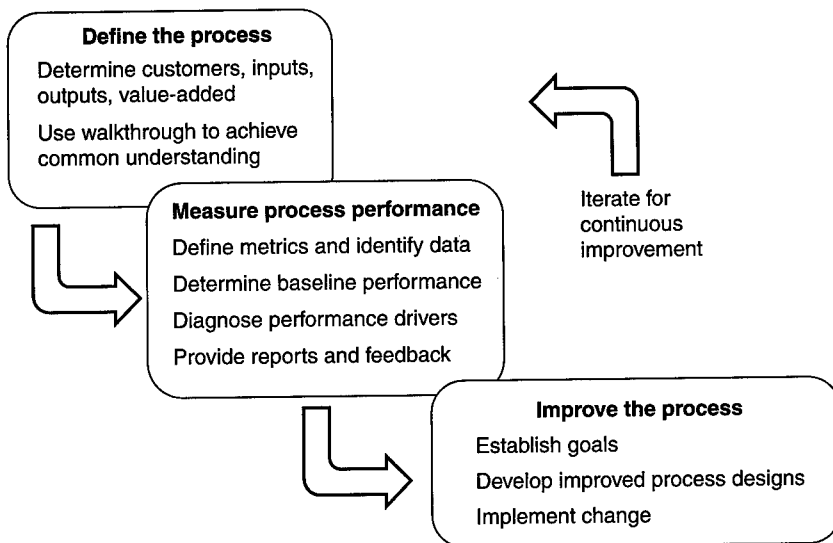


Figure 1.1—The “Define-Measure-Improve” Methodology

performance (Robbins et al., 1998). Another critical area that Marine units have begun attacking is inventory policy and execution.

A more effective set of inventory policies and procedures, backed by metrics, performance reports, and diagnostic tools, could be an important lever for obtaining more effective logistics performance. Figure 1.2 states this point schematically.

This figure conveys the idea that logistics processes are not isolated entities and that improvements in the performance of one can pave the way to superior performance in another (and alternatively, poor performance in one can worsen performance in related processes). In this report, we show that improved order and ship processes can facilitate an improved inventory process by reducing pipeline times and variability, and thus the amount of safety stock needed to cover future demands. Alternatively, improved ways of determining inventory could, with fewer resources invested, reduce order and ship time (because more items will be available in nearby stock locations) and

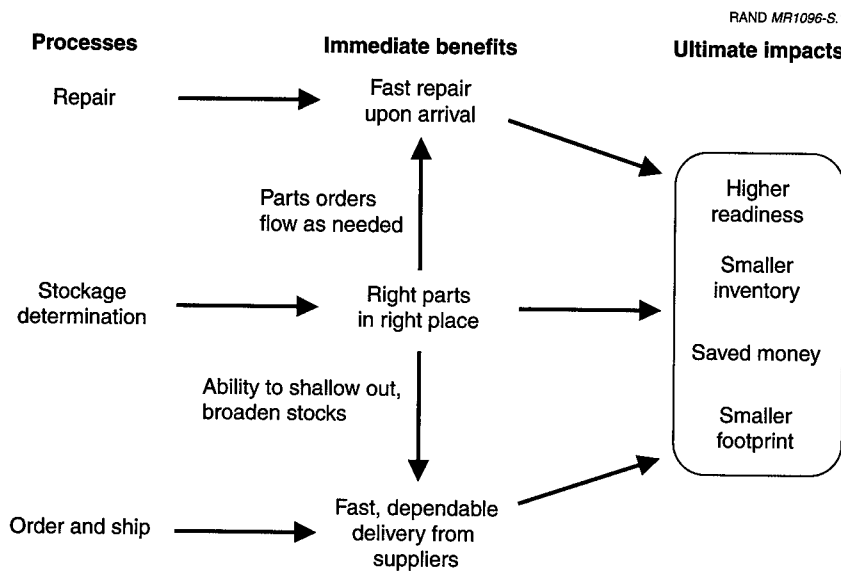


Figure 1.2—Logistics Process Improvements Create Potential for Synergies (Edwards and Eden, 1998)

decrease repair cycle time (because less time will be lost waiting for parts).

In this report we argue that improved inventory management is essential to better logistics performance—hardly arguable or new: the inventory theory literature has made this argument for decades. What is new, however, is that we envision managing the inventory in a very specific organizational and policy framework far different from the abstract insularity in which it is traditionally cloaked. Indeed, while we demonstrate that improved inventory results in better logistics performance, we also show that local inventory is embedded in a larger supply chain and so inventory improvements can be achieved only as part of a larger strategy of improving Marine Corps logistics processes, such as the order and ship process. We offer a new way of looking at inventory, tying it directly into a process of making change (Precision Logistics and the ILC Initiative), linking it to metrics far more meaningful to the customer, and exploring ways in which abstract ideas, algorithms, and methodologies can be

translated into a more satisfied customer and a more capable fighting force.

OUTLINE OF THE REPORT

Chapter Two makes the case for change in inventory policies and practices. It introduces the Marine Corps logistics environment and then presents evidence that recent performance of the supply system in supporting its maintenance customers is in need of improvement. Chapter Three offers a new way of making real change in performance of Marine inventories. In particular, it introduces the "bootstrap" and "dollar-banding" concepts adapted to Marine needs to ensure more bang for the investment buck. In Chapter Four, we demonstrate how our inventory policies support Marine Corps customers more effectively at reduced cost, using a new tool, the "Virtual SMU," and in Chapter Five we discuss how our methodology can be implemented by the Marine Corps *today* and discuss other applications for both the military and civilian sectors. A number of appendices are also included at the end of the report, presenting detailed background information on the Marine Corps supply system and the technical and mathematical details of the new inventory methodology.

THE INVENTORY FULCRUM: MARINE CORPS LOGISTICS POLICY AND PRACTICE

This chapter describes recent logistics performance of the Marine Corps, emphasizing the impact of inventory policies on weapon system maintenance. It offers both illustrative cases and, more formally, some new metrics of overall performance that are at the heart of the "Measure" phase of the Precision Logistics methodology. Before laying out this performance, we first set the stage by describing the relevant parts of Marine Corps logistics.

This report focuses on supply and repair in the active part of the Fleet Marine Force (FMF). The Marine Expeditionary Force (MEF) is the key combat element of the FMF and is composed of three major parts: a ground division, an air wing, and a provider of intermediate logistics (the Force Service Support Group, or FSSG). There are three active MEFs: I MEF at Camp Pendleton, California, II MEF at Camp Lejeune, North Carolina, and III MEF in Okinawa, Japan.

REPAIR AND SUPPLY IN THE FLEET MARINE FORCE

Repair of FMF equipment is performed in any of five echelons of maintenance, four of which reside in the FMF. Echelon 1 is crew-level servicing of vehicles and weapons. This includes standard servicing of the equipment, such as preventive maintenance service checks. Echelon 2 maintenance is performed by mechanics organic to the combat maneuver units, such as infantry battalions. This can include fairly simple repairs, such as removal and replacement of major items. For Echelon 3 and Echelon 4 repairs, the broken equip-

ment is transported back to intermediate repair resident in the MEF maintenance battalion, a part of the FSSG that provides intermediate logistics support in general to the entire MEF. This battalion's five companies perform maintenance for ordnance (from rifles to howitzers), motor transport, electronics, and engineering equipment at the third echelon of maintenance; in Echelon 4 maintenance, the battalion's General Support Maintenance Company does more-demanding work on major assemblies, such as engines and transmissions.

When a broken piece of equipment, whether a principal end item (PEI), such as an Amphibious Assault Vehicle (AAV) or a secondary repairable, such as an AAV transmission, is inducted into the shop, an Equipment Repair Order (ERO) is opened in the Marine Corps Integrated Maintenance Management System (MIMMS). The ERO is a computerized tracking form for all actions on this piece of equipment in a particular shop, including tracking time spent in each phase of repair (inspection, awaiting shop space or parts, etc.), defects noted and man-hours expended, and parts requisitioned and the status of those requisitions.

More complex repairs typically require one or more repair parts to replace those that have failed or are degraded. Each repair action associated with an ERO will have a "layette," or separate parts bin, associated with it, where requisitioned parts are assembled and staged as they arrive. When all parts have been received, the layette's contents are moved to the bench or bay where the broken equipment is located and the repair is effected.

Supply

Each unit performing maintenance may keep small amounts of fast-moving stock nearby (such as bolts and fasteners) in their pre-expended bins (PEBs). For the units' more substantial parts needs, however, they go to the intermediate level of supply for the MEF, maintained by the supply battalion resident in the MEF.

The supply battalion is the main provider of stocks for the entire MEF. Maintaining a suite of warehouses, it divides its inventory into two main parts: consumables, held by the General Account, and repairables, held at the Repairable Issue Point (RIP). These supplies

are maintained and controlled by the SASSY Management Unit, or SMU (SASSY stands for Supported Activities Supply System, the information system used to manage stocks). Thus, while the supply battalion may have many dispersed inventory locations, all are under the single control of the SMU. Units requisition parts by placing orders through their local Asset Tracking for Logistics and Supply System (ATLASS) computer, either through e-mail or with daily submission of computer diskettes to the SMU. The SMU processes all requisitions in batch mode once daily (through an offsite main-frame computer), producing materiel release orders (MROs) the next morning, for picking, packing, and shipping to the requisitioner. Items the SMU does not stock or currently lacks may be sent by the SMU to the wholesale level of supply.

SMU consumable accounts can be replenished from wholesale supplies maintained by the Defense Logistics Agency (DLA) or by providers certified by DLA (e.g., through the Prime Vendor or Virtual Prime Vendor initiatives). Repairable stocks can be replenished by similar sources or from the output of the General Support Maintenance Company.

The dollar value of the on-hand stock in a typical General Account is between \$20 million and \$30 million; the SMU has an annual budget within the same range allocated incrementally on a quarterly basis. The General Account stocks between 10,000 and 20,000 individual types of items ("lines") and the inventory may experience anywhere from 200,000 to 300,000 requisitions a year.

Current USMC guidelines as stated in Marine Corps Order (MCO) P4400.151B allow an item to be stocked at the SMU if there were three demands¹ in the preceding 12 months for items that are combat essential (i.e., have a Combat Essentiality Code [CEC] of 5 or 6), or if there were six demands in 12 months for items that are not combat essential (MCO, 1992). Within these guidelines, the SMUs may add rules for deciding whether to stock an item. The Marine Corps then uses days of supply (DOS) as the primary way to set inventory levels, which is based on multiples of the average daily demand for an item.

¹A demand is defined as a requisition, though a requisition may request multiple quantities of an item.

Inventory levels are typically stated in two parts: ROP (reorder point) and RO (requisitioning objective). See Appendix B for a full explanation of the terms. In accordance with MCO P4400.151B, the ROP is defined as the sum of

- a safety stock factor of either 15 or 30 DOS depending on the item's CEC: 30 DOS for items with a CEC of 5 or 6, indicating the item is combat essential, 15 DOS otherwise, and,
- a variable lead-time factor based on estimated OST for each National Stock Number (NSN) converted to DOS.

Thus, for a hypothetical item with a historical mean demand of three per day, a CEC of 6, and an OST of 11 days, the ROP would be 123 units ($30 \text{ DOS} \times 3 \text{ units/day} + 11 \text{ DOS} \times 3 \text{ units/day}$). The RO for an item is defined as the ROP plus 60 DOS. Thus, for the hypothetical example, the RO would be set at 303 units ($123 \text{ units} + 60 \text{ DOS} \times 3 \text{ units/day}$).

Inventory Performance Metrics

The Marines' usual metric for measuring inventory performance is *fill rate*, which is the ratio of demanded items filled out of inventory to the total number of item demands presented over a given period. This measure of performance can be factored into two other metrics: The *accommodation rate* is the ratio of the number of demands for items locally stocked to the total number of demands, and the *satisfaction rate* is the ratio of the number of demanded items filled to the total number of demands for items locally stocked. That is, the satisfaction rate is a conditional fill rate for only those demands stocked locally, and the overall fill rate is the product of the satisfaction and accommodation rates. Accommodation rate is a measure of inventory breadth, and satisfaction rate is a measure of inventory depth.

These traditional measures capture the supply system's performance *from the perspective of the supply system*. However, we argue that these metrics do not capture the impact of inventory decisions and practice on maintenance customers. To understand that, we need to look first into the detail of how stock affects repair, and then more generally at what new metrics—repair cycle time (RCT) and the ERO fill rate—tell us.

LOOKING INTO THE MIRROR OF CURRENT PERFORMANCE

Why does the Marine Corps need to improve its inventory methodology and execution? Because, to answer bluntly, there is good evidence that current Marine inventory policies and processes are not fiscally efficient, carry too much mass in general, and, most important, do not effectively stock the parts its customers most need.

Although cost is important and will be discussed thoroughly in this report, we believe the primary concern is *support to the warfighter*. Let us turn to that first: how do current Marine retail inventory determination processes support customers in the field?

What is a good measure of supply performance from the perspective of the customer? As we mentioned above, the traditional measure of supply performance is *fill rate*—the percentage of requisitions filled out of local stock rather than being sent to a higher-level wholesale supplier. Fill rates at the best Marine Corps units tend to be fairly high for a military organization, in the 60–70 percent range. Does that mean that the retail suppliers' customers are well-treated and should be satisfied? Not necessarily.

A fill rate of 60 percent means that the customer will be completely satisfied with support from the local supplier 60 percent of the time *if* they only need one part. Most customers are far more demanding than that, because parts needs (especially for repairs) are far more complex. In fact, insofar as repair part demands go, customers need one part about one-third of the time.

Table 2.1 demonstrates this. It looks at recent repair actions in the three active MEFs. It shows the number of requisitions per ERO for all EROs that closed in the month of September 1998 or were still open at the end of the month.

Of those repairs that require parts, more than two-thirds require more than one part. Assuming a fill rate of 60 percent and “random” success in having that part available, a repair requiring two parts has only a 36 percent chance of getting both parts it needs locally ($.60 \times .60$); a repair needing three parts has only a 22 percent chance; four parts reduces that to 13 percent; and so forth. Clearly, the vast majority of repairs are unlikely to get all their parts needs satisfied

Table 2.1
Number of Requisitions per ERO, FMF September 1998
(EROs Using Parts Only)

Number of Requisitions per ERO	Number of EROs	% of EROs	Cumulative % of EROs
1	11,273	25	25
2	6,563	15	40
3-5	10,088	24	65
6-10	7,260	17	82
11-20	4,526	11	93
21-50	2,560	6	99
>50	586	1	100

from local repair in a timely way. They will instead have to go to other sources, principally the wholesale stocks maintained by the DLA in large, dispersed depots. Given lag times in processing requisitions, picking, packing, shipping, and receiving items, such parts will arrive far later than parts satisfied from the local supply shelf. The fact that most—but not all—parts were available locally merely means that the locally available parts will sit waiting in a parts bin dedicated for that repair until the other needed, but unstocked, parts arrive from remote sources.

To get a feel for what this can mean, consider the examples shown in Table 2.2, which uses real-world cases of I MEF repairs as captured in the MIMMS database. It shows the parts required for three repairs, their sources of supply, the order and ship time required to receive the part, and their cost. All parts requisitioned were deadlining (not operational for reason of supply, or NORS = N), i.e., the system could not function until the parts were replaced. Many of the parts were available from nearby sources ("local," meaning typically from bench stocks available at the shop or from the retail source of supply, the General Account [MC1], or the Repairable Issue Point [MC3]). In each case, however, the longest OSTs were associated with parts

Table 2.2
Illustrations of Last Part Received

ERO	Equipment	Repair Part NSN	Source of Supply	NORS	PRI	OST	Parts Cost
ID152	M998 CARGO	2530013579708	local	N	06	2	208.17
		2530013579776	local	N	06	4	185.30
		4710013581943	S9C	N	06	7	8.27
		4710013579968	S9C	N	06	9	3.40
M4460	ISOBED TRU	2815001780268	MC3	N	06	0	13,833.00
		2520009715016	MC3	N	06	1	4,366.00
		2540003044306	local	N	06	1	320.11
		2930009192875	local	N	06	2	11.50
		3110005542975	local	N	06	5	0.00
		2815004042916	local	N	06	6	663.90
		4730004213924	S9C	N	06	8	6.44
		4730010955833	S9C	N	06	8	2.48
		4730003652690	S9C	N	06	15	3.92
		5340001343380	S9I	N	06	15	4.14
		2815011867251	MC3	N	03	0	28,091.63
		5305012157281	local	N	03	1	7.36
		5945011451878	local	N	03	1	16.52
		5310006379541	S9I	N	03	7	11.04
M6E59	MK48 PWR U	5310007851762	S9I	N	03	7	33.68
		4730001881864	S9C	N	03	8	0.27
		4730002782719	S9C	N	03	8	1.42

NOTE: NORS = Not operational for reasons of Supply. PRI = priority.

from a wholesale supply source.² In the first repair, an extra five days is tacked onto repair time to await parts not available locally; in the second repair, another nine days, and in the third, an additional five.

In general, not having all parts available locally leads to increased RCTs, as Figure 2.1 illustrates. The figure shows repair times³ for each of the three MEFs for May through December of 1997. For each MEF, we show the spread of repair times (reading from left to right) if one or more high-priority parts (priority 06 or higher) must come from a wholesale source, if all high-priority parts are available immediately from local supply, and if no high-priority parts are

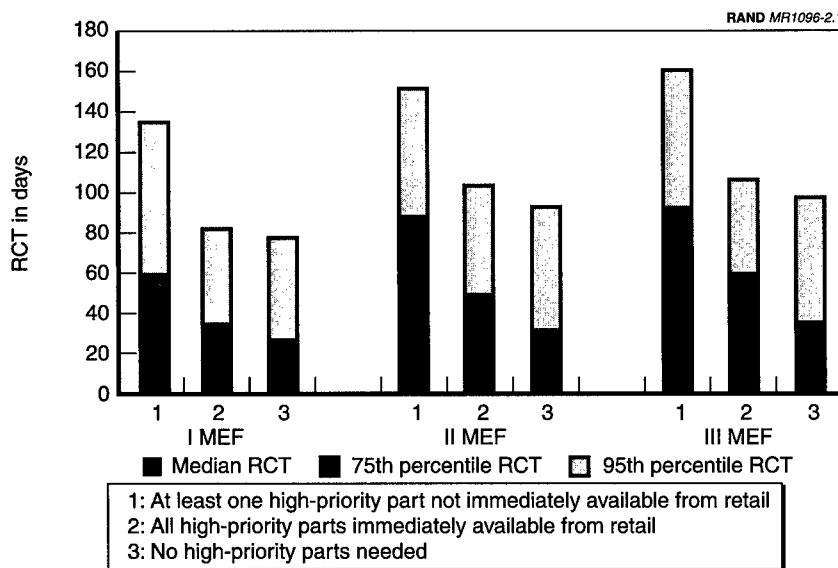


Figure 2.1—Impact on RCT of Parts Supplied from Wholesale

²The two wholesale supply sources in these examples were two inventory control points (ICPs), which are part of the DLA: S9C (Defense Supply Center, Columbus, Ohio) and S9I (Defense Industrial Supply Center, Philadelphia, Pennsylvania).

³Measured in percentiles: 50th, 75th, and 95th percentiles of RCT.

required for the repair at all. Of special interest is the difference between the middle bars (2) and those on the left (1). Needing *any* part at all from wholesale results in a 30 percent performance penalty in terms of repair cycle time on average.

Looking at Table 2.2 again, in each case, a part supplied from wholesale held up the repair for some number of days until it was shipped and received. Look, however, at the price of the parts requisitioned: \$3.40, \$4.14, and \$.35. The other parts required, which were acquired much more quickly, tended to be more expensive; indeed, two parts costing more than \$13,000 and \$28,000 were made available the day they were ordered. It is not an infrequent occurrence that repairs may be held up a significant time awaiting the arrival of what may be a very cheap part.

Table 2.3 presents cases where even a single cheap part can lengthen repair time considerably. The table shows 11 examples of Marine PEI repairs, the repair time, and the “short parts time” (time awaiting repair parts) for each repair. It then shows the order and ship time and cost for the last repair part needed to complete the repair (the part identified by its NSN). To put these cases in larger context, we next turn to RCT, and the drivers of Marine Corps repair performance.

RCT Performance in the FMF

The Marine Corps uses a comprehensive definition of RCT (especially for PEIs), tracking repair time from the breakdown of the system through all stages of repair and movement among echelons of repair until the system is fully capable again: we call this *the repair cycle time for critical repair actions*. Appendix A gives details on how this metric is defined and calculated.

Tables 2.4, 2.5, 2.6, and 2.7 give some evidence of recent RCT performance for PEIs in the three active MEFs during calendar year 1998, broken out for PEIs and secondary repairables. The first part of the table for both PEIs and secondary repairables shows RCT performance, listing both average time and the percentiles that

Table 2.3
Examples of Inexpensive Parts Driving RCT

ERO	PEI in Repair	Repair Time	Repair Short Parts Time	Last Repair Part OST	Cost	Last Repair Part NSN
S2L52	M998 CARGO	69	61	56	7.50	5330011748145
MP654	M813A1 TRK	44	32	32	5.22	3110005542975
S2L66	M998 CARGO	50	46	30	17.65	3110001005920
UQ751	40 TN TRL	49	39	28	2.44	4730012649488
AYR41	M813A1 TRK	84	28	25	0.42	5305000881302
AS037	M998 CARGO	55	25	22	5.38	5305005434372
M4595	M813A1 TRK	42	20	20	0.80	5310011552503
PG036	M998 CARGO	29	26	22	13.37	3030012826968
PG040	TRUCK M923	28	25	22	14.13	3030011181318
TC221	M1038 CARGO	39	33	22	14.44	5977011939931
PML04	M998 CARGO	36	36	21	2.75	5330011940473

describe the distribution of repair times. The bottom half of the table gives evidence of where time is spent in the repair cycle, based on status code histories entered into the MIMMS periodically throughout the repair. The table shows that awaiting parts time ("short parts") is a major contributor to both PEI and secondary repairables repair times, consuming close to half of the total RCT.⁴

⁴Status times are reported for all repairs, whether they need parts or not. As was discussed in Chapter One, for just those repairs that require parts, awaiting parts time thoroughly dominates the entire repair time.

Table 2.4
Principal End Item Repair Cycle Time—CY 1998

	I MEF	II MEF	III MEF
Number	17,560	11,528	6,748
Average	37	51	55
50th	24	30	33
75th	48	65	67
95th	119	164	183

Table 2.5
Principal End Item Average Time in
Repair Status—CY 1998

	I MEF	II MEF	III MEF
To first status	2	3	5
Inspection progresses	3	4	4
Short parts	12	18	19
Short technician	1	2	1
Repair progresses	2	2	2
Awaiting evacuation	1	2	3
Evacuated to higher echelon	8	13	11
Pickup from higher echelon	1	1	1
All others	6	5	8
Repair completed	1	1	1

Table 2.6
Secondary Repairable Repair Cycle Time—CY 1998

	I MEF	II MEF	III MEF
Number	20,742	15,154	6,755
Average	33	44	44
50th	22	21	24
75th	43	49	54
95th	103	166	149

Table 2.7
Secondary Repairable Average Time in
Repair Status—CY 1998

	I MEF	II MEF	III MEF
To first status	2	2	4
Inspection progresses	2	7	4
Short parts	8	16	16
Short technician	3	1	6
Repair progresses	1	3	2
Awaiting evacuation	2	1	0
Evacuated to higher echelon	8	2	0
All others	5	9	9
Repair completed	2	3	3

The ERO Fill Rate: Definition and Performance

Short parts time is driven by many factors: ordering and receiving processes, responsiveness of the wholesale system, criteria used by the wholesale system for acquiring and stocking parts, and other factors. But a critical factor, because it is such a major and nearby source of supply, is the local stock held by the SMU. To the extent that the SMU stocks and has available critical parts needed for a repair, short parts time will be lessened, RCT will decrease, and the availability and readiness of weapon systems will thereby increase. However, as we showed before, having *most* of the critical parts needed for a repair typically does little good—repairs cannot be completed until *all* the critical parts are available. This way of looking at the performance of the local supply system is the heart of a new Marine Corps metric, the “ERO fill rate.” The ERO fill rate shows how often local supply is able to immediately provide *all* the high-priority parts needed to complete a “critical” repair. It is an all-or-nothing score: if all needed parts are available on the shelf when ordered, the ERO fill rate score is “one”; if even one part must be obtained from sources outside the base, even if 99 other parts also needed for the repair are immediately delivered from SMU stocks, the ERO fill rate score is “zero.” (See Appendix A for details on how the ERO fill rate is calculated.) That is, the ERO fill rate is a measure of the local supply system’s performance *from the perspective of the mechanic*.

ERO fill rate performance for each of the active MEFs in calendar year 1998 is given in Table 2.8. It shows that less than half of all PEI repairs that require high-priority parts get all of those parts from available stocks in the local warehouse—the rest must pay a “wholesale” penalty for the remainder of the needed parts.⁵ For secondary repairables, the SMUs are even less able to provide effective support. Thus, a more effective Marine logistics system—providing fast repair and yielding higher weapon system readiness—is missing one critical linchpin: a supply system structured to meet the *real* needs of the weapon system maintainers. To enhance that performance, a new approach to inventory is required. The next two chapters present that new approach.

Table 2.8

**ERO Fill Rate Performance by MEF (CY 1998): Overall,
by PEIs, and by Secondary Repairables**

MEF	Aggregate		PEIs		Secondary Repairables	
	Number	%	Number	%	Number	%
I	9,519	49	4,976	56	4,543	41
II	7,452	38	3,707	49	3,745	27
III	3,870	36	1,858	46	2,012	27

⁵This delay penalty is not associated with or a result of any particular supplier, such as DLA. Rather, the delay is often a result of inefficiencies in processing the paperwork to order a part, compared with supplying it from local stock. Much of the delay is within the control of the SMU and can be minimized through process improvement. Nonetheless, even the most efficient process for ordering and supplying nonlocal stocks will generally involve some delay as compared with the supply of local stock, incurring at least additional shipping times.

Chapter Three

A LEVER FOR RAISING PERFORMANCE: AN IMPROVED INVENTORY METHODOLOGY

One almost universal function of any military supply system is to warehouse goods in anticipation of customer demands. Similar requirements exist in portions of the commercial sector and other operations where the advantages of holding stock locally offset the cost of maintaining and managing the stock. An advantage of maintaining local stock is the ability to satisfy a customer demand almost immediately. In particular, local stock is often used as a buffer between the customer and the vagaries of the rest of the supply system. Maintaining local stocks requires determining which items to stock and in what quantities. In peacetime, overstocking ties up funds in assets that turn over very slowly; in wartime it results in masses of unused materiel that must be transported or left behind. Understocking defeats the purpose of maintaining local inventory when items are not available and can have a large impact on wartime sustainability.

This chapter presents new techniques to improve the Marine Corps retail inventory performance, including ERO fill rate, through changes in supply policy and methodology. It concentrates on two particular techniques: properly calculating ("bootstrapping") the reorder point and broadening ("dollar banding") the inventory. These techniques are the "levers" through which substantial improvement in inventory performance is possible. Each technique is discussed in a separate section. The chapter assumes a basic understanding of inventory theory. Such a background is necessary to follow the terminology and development of the concepts. Appendix B provides a brief introduction for those readers unfamiliar

with the subject.¹ Two terms discussed in the appendix are critical here: reorder point (ROP) and requisitioning objective (RO). The ROP is the stock level at which a replenishment order is made and represents the amount of stock held to cover future demands until the replenishment is received. The RO is the maximum on-hand quantity.

In the next section, we begin by describing a method to set the ROP by risk of stock-out. This is a significant departure from the current method of doing business in the Marine Corps, in which standardized rules are applied across all items, resulting in varying service levels *not under the inventory manager's control*. Such a system denies the inventory manager the most basic and natural management criteria: setting the inventory level of each item according to the level of service that the inventory manager feels is most appropriate. After that, we present the idea of "dollar banding," which incorporates both cost and demand in the stock decision, as a new way for Marine Corps managers to think about how to stock. In its simplest form, dollar banding promotes higher service levels for less-expensive items, all other things being equal, and often results in a broadened inventory.

Taken together, these two techniques prescribe new ways for Marine Corps logisticians to think about how to set inventory. Of course, all this will ultimately be tied back to how these new inventory methodologies contribute to improved ERO fill rates, from which significant improvements in the standard inventory metrics also follow. This chapter concludes with a discussion of the impact of the order and ship process on new inventory strategies like these. We argue that a better performing inventory *must* be accompanied by improvements in the SMU's order and ship process. Ultimately, logistics processes are entwined with each other, and real improvement demands an attack on all fronts.

¹Also, in Appendix C we discuss some common inventory misconceptions, some of which are codified in Marine Corps Orders.

ILLUSTRATING PROBLEMS WITH CURRENT MARINE CORPS INVENTORY METHODOLOGY

There are a number of weaknesses with the Marine Corps methodology. First and most important, it eliminates information on demand variability contained in the data by using only averages. For example, an item ("Item 1") that deterministically has one demand for one unit a day for 365 days has exactly the same DOS (average daily demand), and thus the same RO and ROP, as another item ("Item 2") that has only three demands for, say, 100, 125, and 140 units on three separate days and no demands in the other 362 days. Indeed, consider how the inventory position and on-hand stocks might look over the course of a year under the following realistic, simplifying assumptions, which result in both items having an ROP of 40 and an RO of 100.²

- The items both have CEC codes of 5.
- The OST for both items is 10 days.
- Thus, the operating level would be set at 60 DOS and the lead time to 10 DOS for both items.
- So, under current USMC guidance, both items would have equal safety stock, computed as $30 \text{ DOS} \times 365 \text{ units ordered in the past year} / 365 \text{ days} = 30$.

Figures 3.1 and 3.2 show that the ROPs would have worked out very poorly in practice in the next year if the two items saw demand pat-

²As discussed in the previous chapter, inventory levels are typically stated in two parts: ROP (or reorder point) and RO (requisitioning objective). See Appendix B for a full explanation of the terms. In accordance with MCO P4400.151B, the ROP is defined as the sum of

- a safety stock factor of either 15 or 30 DOS, depending on the item's CEC: 30 DOS for items with CEC of 5 or 6, indicating the item is combat essential, 15 DOS otherwise, and,
- a variable lead-time factor based on estimated OST for each NSN converted to DOS.

Thus, for a hypothetical item with a historical mean demand of three per day, a CEC of 6, and an OST of 11 days, the ROP would be 123 units ($30 \text{ DOS} \times 3 \text{ units/day} + 11 \text{ DOS} \times 3 \text{ units/day}$). The RO for an item is defined as the ROP plus 60 DOS. Thus, for the hypothetical example, the RO would be set at 303 units ($123 \text{ units} + 60 \text{ DOS} \times 3 \text{ units/day}$).

terns similar to those used to calculate the RO and ROPs. That is, these figures show that neither item is well supported using the current methodology. In the case of Item 1, the safety stock is unnecessary but the fixed DOS methodology results in 30 units kept in stock to guard against fluctuations that never occur. If future demand remains at exactly one demand a day then these 30 units are destined to sit on the shelf and never be used.³ For Item 2, demand occurs in large batches, and the result is that either the inventory is full up to the RO with items not being demanded or, when they are demanded, there is insufficient inventory to fill the entire order. Batches are a commonly observed phenomenon in Marine Corps data.

PROPERLY CALCULATING THE ROP: THE BOOTSTRAP ROP

As is discussed in Appendix B, "Basic Inventory Theory and Terminology," the ROP should be set based on choosing the risk of stock-

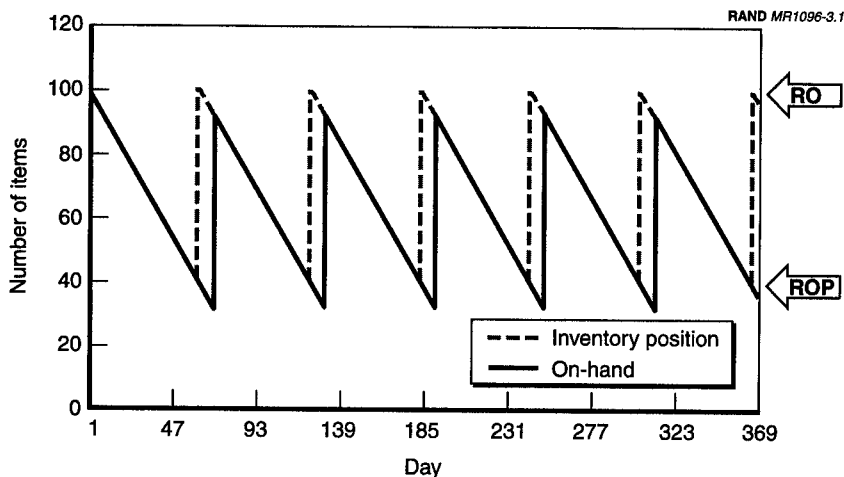


Figure 3.1—The Performance of Hypothetical Item 1 Using the Current Marine Corps Methodology for Setting the RO and ROPs

³Of course, such perfectly deterministic patterns are not likely to occur often in practice, but similar results could well occur in items that have low variability in demand.

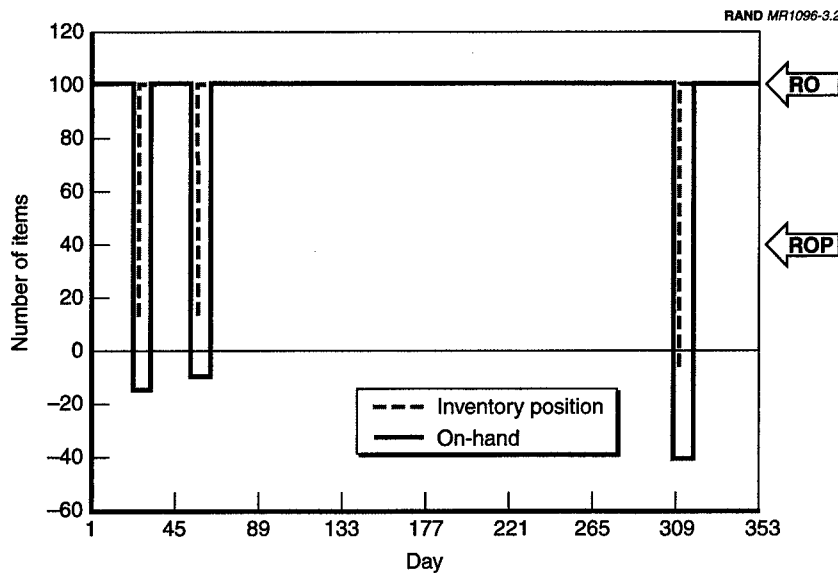


Figure 3.2—The Performance of Hypothetical Item 2 Using the Current Marine Corps Methodology for Setting the RO and ROPs

out when the on-hand stock is at the ROP. The current methodology based on the average demand (using the DOS methodology) inherently cannot do this. By using the average demand, the DOS methodology fails to account for unavoidable variability that exists throughout the inventory process. This variability enters into the process through

- fluctuations in demand;
- variation in order quantities; and,
- irregularities in lead times, in particular wholesale order and ship times.

How the Bootstrap ROP Methodology Works

As an alternative to the existing methodology, we created a new computational algorithm for setting ROPs, based on the “bootstrap”

technique first developed by Brad Efron of Stanford University (see Efron, 1979, and Efron and Tibshirani, 1993, for example).⁴ We call our method the *bootstrap ROP*. Applying the bootstrap was first proposed by Bookbinder and Lordahl (1989) with subsequent work by Wang and Rao (1992). Our approach is motivated by the shortcomings of the current DOS methodology, as previously discussed and illustrated, and the difficulties in applying existing parametric techniques⁵ to Marine Corps data.⁶

Appendix D provides the technical details of the bootstrap methodology. Here we will try to give some intuition into how and why the method works. To start, imagine a simple inventory problem where we know the lead time is exactly 10 days, and we have a lot of historical data. Then, to find the risk of stock-out for setting the ROP to some number x , one could simply extract from the data all the demands for each set of 10 days after the inventory position reached the ROP and count the fraction of times out of the total number in which total number of demands exceeded x . This fraction would be a good estimate of the risk of the chance of a stock-out once the inventory position got down to x . Essentially, this is what the bootstrap procedure does.

The basic idea is to construct a distribution of demands that occurred in the lead time after a replenishment buy order was placed (commonly called the distribution of lead-time demand, or LTD, in inventory theory literature). The difficulty with existing historical data is that they are generally insufficient to construct such a simple empirical distribution as in the previous paragraph. For example, one year of Marine Corps' historical data with average lead times of

⁴"The use of the term bootstrap was derived from the phrase *to pull oneself up by one's bootstrap*, widely thought to be based on one of the eighteenth century *Adventures of Baron Munchausen*, by Rudolf Erich Raspe. (The Baron had fallen to the bottom of a deep lake. Just when it looked like all was lost, he thought to pick himself up by his own bootstraps.)" (Efron and Tibshirani, 1993). The name for the statistical procedure is apt because it tends to work when the data are insufficient for other techniques.

⁵A "parametric technique" assumes a particular probabilistic form for the data. For example, one technique assumes that the number of demands in a particular lead time period is normal.

⁶See Fricker and Goodhart (forthcoming) for a comparison of the bootstrap procedure to parametric alternatives.

20 to 30 days makes for too few data points. The bootstrap methodology overcomes this difficulty in a sophisticated way, using the limited data multiple times. As Appendix D discusses, the data must meet some assumptions for the result to be accurate; users should be aware of these and adjust the methodology to account for violations in the assumptions.

Returning to the hypothetical Item 1 (Figure 3.1), the bootstrap methodology would set the ROP at 10 if the inventory manager specified a zero risk of stock-out. This, of course, is obvious from the knowledge that one demand occurs every day and the lead time is exactly 10 days. However, it is comforting confirmation that the methodology gives exactly the expected answer. Item 2 (Figure 3.2) is more complicated, and it is in these more complicated cases that the power of the bootstrap becomes apparent. For Item 2, which has very lumpy demand, the bootstrap sets the ROP to zero for any risk greater than 7 percent; otherwise it sets the ROP above zero in increments according to risk. For example, a specified risk between about 5 and 7 percent would give an ROP of 100, and a risk between about 2 and 5 percent would give an ROP of 125. Looking at the plot in Figure 3.2, this is a sensible thing to do. If the inventory manager can accommodate a high level of risk⁷ (in this case, running out of

⁷Note that the risk is calculated under the assumption that the on-hand quantity is at the ROP, so that the full ROP number of items are available at the start of the lead time. This condition may not actually occur in practice, however, particularly when orders occur in large batches (as with Item 2). In such a case, the inventory position may overshoot the ROP, leaving the on-hand quantity somewhat less than the ROP when the replenishment order is placed, resulting in a greater actual risk of stock-out than the specified risk.

This "overshoot effect" lessens as the size of the batches decreases, vanishing with items that are ordered singly or in fixed lot sizes. It increases as order quantity increases and order frequency decreases. Item 2 is an extreme case of such behavior, and the worst situation for the bootstrap procedure is realized when the RO is set only slightly above the ROP. Then, even with the on-hand quantity at the RO, when a requisition is received for 100 to 140 items, virtually no stock is left to cover during the lead time.

In this case, we can bound how far the actual risk differs from the desired risk, because the net result is that any requisition arriving in the lead time period results in stock-out. This is equivalent to setting the ROP to zero, which gives a 7 percent risk of stock-out (that is, a 7 percent chance that a requisition will arrive during the lead time). Thus, the most the actual risk can differ from the specified risk is 7 percent. For higher risk levels in this example, and for items less extreme than the example, the difference between actual risk and specified risk should be even smaller.

stock more than 7 percent of the time⁸), then it is reasonable not to hold any safety stock because it is ordered only infrequently and it is relatively unlikely that another order will come in within 10 days of the inventory position getting to zero. However, if the inventory manager is more risk-averse, the ROP must be set to 100 or greater, depending on the degree of risk-aversion. Setting the ROP at 100 or greater in this case is sensible because it guards against the large batches. Setting the ROP between zero and 100 will likely result only in a partially filled order should a demand occur in the lead time.

The bootstrap methodology can be applied in a number of ways. The most realistic, in the sense of incorporating both demand and lead-time variability in the risk assessment, is to use actual historical demand data and actual OST data in the calculations. The result is an estimate of the complete lead-time demand distribution, from which ROPs can be selected that account for the overall risk of stock-out. Alternative OST distributions can also be used in the calculations to evaluate how the inventory position would change, for example, if OSTs were reduced. In much of the work to follow, we fixed the OSTs at 18 days to examine just this possibility and to illustrate the benefits of improved OSTs. Once the bootstrap methodology is applied, the inventory manager must specify the desired risk for each item. This would be a daunting task on an item-by-item basis; our solution will be to apply "dollar banding" to specify risks for sets of items grouped according to cost.

Application to USMC Data

To demonstrate the methodology, we applied it to demand data for February 1997 through July 1998⁹ from the I MEF and generated

This technical distinction pales in significance when the bootstrap method is compared with the USMC methodology. Using the bootstrap on Item 2, one might specify a 3 percent or 5 percent risk and actually experience a 7 percent risk (under the assumption that future demands look like past demands). Yet, under the current USMC methodology and the same assumptions, one would experience a 100 percent risk—every order results in a stock-out.

⁸While a "7 percent rule" might not seem all that great, the risk would be calculated and applied individually to the thousands of parts stocked.

⁹The demands consisted of all ZOA transactions augmented with 1st FSSG *direct ordering system* transactions. A ZOA record was required in order to specify the date

alternative ROPs (which we call the "bootstrap ROPs"). We compare these results with those of the April 27, 1998, General Account Balance File (GABF), which lists by NSN the actual RO, ROP, on-hand, due-in, and due-out stocks for the I MEF General Account.¹⁰

Using the bootstrap methodology we calculated ROs and ROPs that reflect how the I MEF General Account might be stocked assuming various risks between 0.05 and 0.01.¹¹ We used a variable lead time in the bootstrap ROP calculations derived from actual (non-back-ordered) replenishment lead times of 30 days or less.¹² For the sake of simplicity, the RO was set using an operating level based on days of supply and the on-hand quantity was set randomly between the RO and the ROP; the details are discussed in Appendix D.

Table 3.1 compares the value of an inventory based on these new RO and ROPs with the actual value of the April 27 GABF. In the third col-

Table 3.1
Comparison of the Actual April 27 GABF to Various
Alternative Bootstrap GABFs

	Actual GABF	Bootstrap GABFs			
		$r = 0.05$ OL = 30 DOS	$r = 0.03$ OL = 30 DOS	$r = 0.01$ OL = 30 DOS	$r = 0.01$ OL = 90 DOS
Number of lines (NSNs) stocked	13,159	14,014	17,682	32,645	32,645
Value at the ROP	\$13.5	\$4.7	\$6.7	\$13.2	\$13.2
Value on-hand	\$25.5	\$5.1	\$7.2	\$13.7	\$21.3
Value at the RO	\$23.6	\$6.4	\$8.8	\$16.9	\$24.5

NOTE: Dollar figures are expressed in millions, r is the risk, and OL is the operating level in days of supply.

when the demand was observed at the SMU. It should be noted, however, that some transactions do not generate a ZOA and are thus not captured in our demand stream.

¹⁰ Note that the RIP is *not* included in these calculations.

¹¹ A risk of 0.01 means that one time out of a hundred we expect the number of demands in the lead time to exceed the ROP and thus cause a stock-out.

¹² We have, thus, assumed that the current OST process improvement effort will continue and eliminate excessively long replenishment times, because about 10 percent of the non-back-ordered replenishment times were greater than 30 days.

umn, the table shows the results from the bootstrap GABF with a risk (r) equal to 0.05 and an operating level (OL) of 30 DOS. The results show that about 1,000 more lines are stocked compared with the actual GABF, with a value at the ROP of less than half the actual GABF's value at the ROP. The other bootstrap GABF options illustrate how the number of lines and the dollar value of the RO and ROP inventory increase for decreasing levels of risk.

Table 3.1 shows that the SMU can stock *many* more lines at no greater cost, indeed, possibly at much less, than the actual GABF. And, as Chapter Four will show, a bootstrap GABF based on a risk of 0.05 and an operating level of 30 DOS *performs as well as the actual GABF in terms of fill rate, satisfaction rate, accommodation rate, and ERO fill rate—at less than one-half the inventory cost*. The bootstrap GABF based on a risk of 0.01 and an operating level of 30 DOS *outperforms the actual GABF, achieving ERO fill rates of more than 70 percent and even higher fill rate metrics at three-fourths of the inventory value at the RO*.

Examples of Bootstrap ROPs

To demonstrate how the bootstrap ROP works, it is instructive to look at some specific cases. Table 3.2 shows six NSNs with varying demand patterns along with three possible bootstrap ROPs corresponding to risks of 0.01, 0.05, and 0.10 for a fixed 18-day lead time (for illustration). The first three cases had a total of five requisitions for the year from February 1997 to January 1998, with varying total quantities requested. The last three cases have a larger number of requisitions (from 50 to 104).

Case 1 shows an NSN that had a total of five requisitions in five different months, with the greatest number (50) in April, and a total quantity of 110 items ordered for the 12-month period. For the three different risk levels considered, $r = 0.10$, 0.05, and 0.01, the bootstrap ROPs are 15, 25, and 60, respectively. That is, as the allowable risk is decreased (equivalently, the service level is increased), the ROP is increased. At $r = 0.01$, the lowest level of risk, the ROP is set greater than the largest requisition, allowing for the possibility that both a requisition for 50 units and a requisition for 10 units (or perhaps

Table 3.2
Various NSNs and Their Annual Demand Pattern for One Year with Three Possible
Bootstrap ROPs Corresponding to Risks of 0.01, 0.05, and 0.10

Case	NSN	Total Annual			Bootstrap ROP			1997												1998											
		Unit Price (\$)	Quan- tity ^a	Num- ber ^b	$r = 0.10$	$r = 0.05$	$r = 0.01$	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan												
1	7641014005223	0.01	110	5	15	25	60	0	15	50	10	0	0	0	0	0	10	25	0												
2	8415010333520	0.06	488	5	50	108	282	0	0	282	0	60	48	98	0	0	0	0	0												
3	2940010882429	1,652.00	5	5	0	0	5	0	0	0	5	0	0	0	0	0	0	0	0												
4	1005007266131	808.00	63	50	9	13	18	0	26	3	5	1	1	2	2	1	5	13	4												
5	8465009651705	1.55	3,095	100	316	395	567	526	355	144	216	641	252	54	40	630	5	128	104												
6	5325011917555	4.57	174	104	12	24	74	7	13	24	6	25	5	0	2	0	3	77	12												

^aQuantity requested.

^bNumber of requisitions.

some other combination of requisitions) arrive in the lead time. This example shows how an inventory manager can specify a risk level and the methodology will set the ROP appropriately.

Case 2 is roughly equivalent to Case 1, though the quantities requested on each of the five requisitions are larger. Here the ROPs for the various risk levels are set larger, with $r = 0.01$ set to the maximum observed in April, 282 units. Case 3 shows a different situation in which five expensive items (\$1,652.00), each on a different requisition, all occurred in a batch in a short period of time. In this case, the bootstrap ROPs for the higher risk levels set the ROP to zero, but for the smallest risk (highest service level) it is set to five units. Here the methodology also performs according to intuition: it either sets the ROP to cover the entire batch of five or not to cover them at all, depending on the risk that the inventory manager is willing to assume. A comparison of Cases 1 and 2 with Case 3 hints at how dollar banding will later be incorporated into the methodology: for *many* inexpensive items, such as those in Case 2, a low risk can be achieved (moving from $r = 0.05$ to $r = 0.01$ for Case 2 would cost only \$10.44 and for Case 1 only \$.35) for the cost of assuming slightly more risk for just one NSN. That is, in Case 3 it would save \$8,260 to allow $r = 0.05$ instead of $r = 0.01$, which would, in turn, pay for many cheap items to have a high service level.¹³

Cases 4 through 6 show NSNs with higher requisition activity, their demand patterns, and the recommended bootstrap ROPs for the three risk levels. Case 4 shows that the largest bootstrap ROP is smaller than the greatest month (March) because the requisitions during that month were sufficiently spread out. Case 5 is similar, though with even higher requisition activity and larger quantities ordered. Case 6 has a high number of requisitions but a low number of units per requisition. Taken together, the six cases in Table 3.2 cover a wide spectrum of demand patterns and unit prices. All of these cases suggest that the bootstrap methodology sets reasonable ROPs that correspond with intuition when reviewing the monthly demand patterns.

¹³This result nicely complements and reinforces our maintenance philosophy as well: end items should never be held down awaiting an inexpensive part from wholesale.

INCORPORATING "BANG FOR THE BUCK": DOLLAR BANDING

Dollar banding incorporates the price of an item in the inventory decision and allows the breadth of the inventory to be significantly increased for the same inventory investment.¹⁴ Essentially, the less expensive an item, the more liberal the inventory decision should be, leading to higher RO/ROP quantities. The idea is that stocking extra quantities of cheaper items is inexpensive insurance against surges in demand and the other types of variation inherent in the supply system. Dollar banding is often applied to great advantage by assuming a slightly higher risk of stock-out for a few low-demand, expensive items and using the savings to achieve significantly higher service levels for many inexpensive items.

Simply stated, dollar banding makes inventory decisions a function of *both demand and unit price* rather than demand alone. The concept is commonly used in commercial industries, such as retail stores for vehicle repair parts, where common, inexpensive parts are routinely stocked while more expensive, less frequently demanded parts are held at a warehouse rather than in stock.

A Simple Example

Here we present some simple simulations to demonstrate the effect of changing the inventory rules to incorporate cost. These simulations use actual demand data from I MEF from February 1, 1997, to January 31, 1998, in determining what to stock. For each simulation, the ERO fill rate for February–April 1998 is calculated and the performance of the rule is compared in terms of ERO fill rate, number of lines stocked, and cost of the inventory at the ROP.

The simulations were conducted as follows: (1) a stock criterion was first defined, such as the three demands in 12 months, (2) the stock criteria were applied to actual I MEF data from February 1997 through January 1998 to generate a list of "stocked" NSNs, and (3)

¹⁴Our RAND colleague, Dr. Ken Girardini, first suggested the idea of dollar banding to us, after he demonstrated significant inventory performance improvement using the technique in the Velocity Management program.

finally the NSNs were *assumed* to be perfectly stocked for the next quarter. Thus, if a “stocked” NSN is demanded in the new quarter it is always filled out of local stock (i.e., no stock-outs exist in this perfect world) and, conversely, if the NSN is not “stocked” then it is never filled out of local stock.

As a baseline, consider the MCO P4400.151B stocking rules applied without regard to financial or other constraints. That is, CEC 5 and 6 items with three or more requisitions in the year (February 1997 through January 1998) and all other CEC items with six or more requisitions are stocked to such a depth that they are always available for issue. For I MEF this results in an ERO fill rate of 76.3 percent for the February through April 1998 “quarter.” Using these criteria, I MEF would have stocked 10,665 lines.¹⁵

This ERO fill rate figure is, of course, higher than what was really experienced in I MEF because: (1) the SMU is not financially able to stock all items that meet the demand-based stocking criteria, (2) stocked items experience back-orders because of stock-outs at the wholesale level, and (3) as discussed in the previous subsection, the actual ROPs are not correctly calculated, resulting in retail stock-outs. However, this example provides a “perfect world” baseline from which the effect of stocking rules based on both demand and cost can be evaluated.

So, under dollar banding, one should change the stocking rules to account for unit cost in the stock decision, particularly for inexpensive items. The idea is that the less expensive an item is, the more forgiving, in terms of number of demands, the stock decision should be. Consider four possible rules:

- The demand-only or baseline rule uses the MCO P4400.151B stocking rules, which do not consider cost. This is the baseline rule, described above, against which the following three dollar-banding rules will be compared.
- The simple demand and cost rule is to stock an item if three or more demands occurred in the past year and the item had a unit

¹⁵We are using a slightly decreased demand stream than what actually occurred, which is why the number of lines stocked here according to actual rules is slightly decreased from the number shown in the actual GABF in Table 3.1.

price of less than \$50 or if the item had six or more demands and any unit price.

- The medium demand and cost rule graduates this a bit more finely. Stock an item if it has
 - one or more demands and costs less than \$10;
 - two or more demands and costs less than \$25;
 - three or more demands and costs less than \$50; or,
 - six or more demands at any unit price.
- The complicated demand and cost rule provides additional restrictions on expensive items while relaxing restrictions on the cheaper items. Stock an item if it has
 - one or more demands and costs less than \$10;
 - two or more demands and costs less than \$50;
 - four or more demands and costs less than \$100;
 - six or more demands and costs less than \$500;
 - eight or more demands and costs less than \$2,500; or
 - 10 or more demands at any unit price.

Table 3.3 shows the expected performance of the three dollar-banded stocking rules versus the MCO P4400.151B demand-only rule. The simple cost and demand rule slightly improves the ERO fill rate performance by adding 1,735 lines while saving almost \$1 million in inventory cost at the ROP. The medium cost and demand rule further improves on the inventory performance by adding more than 10,000 additional lines with a negligible increase in cost. Finally, the complicated cost and demand rule again projects improved inventory performance while decreasing the inventory cost to the lowest of the four options. That is, the complicated rule gives an expected 9.2 percent increase in simulated ERO fill rate while saving \$1 million dollars in inventory investment. This example illustrates that dollar banding can improve inventory performance while decreasing inventory investment.

Table 3.3
Comparison of the Various Dollar-Banding Schemes with the
Demand-Only Baseline That Follows USMC Guidelines

Scheme	Simulated ERO Fill Rate (%)	Number of Lines	Inventory Cost (at the ROP)
Demand-only (baseline)	67.1	10,665	\$10.85
Simple cost and demand	67.5	12,463	\$9.86
Medium cost and demand	75.3	22,913	\$9.90
Complicated cost and demand	76.3	23,675	\$9.82

NOTE: Dollar figures are expressed in millions.

Combining the Dollar-Banding and Bootstrap ROP Methodologies

In the previous simple example, we used dollar banding to determine whether an item would be stocked. However, the application of dollar banding would proceed differently in combination with the bootstrap methodology. As was mentioned in the previous section, dollar banding can be used to specify the bootstrap ROP service level. In particular, expensive items should be assigned a lower service level and inexpensive items should be assigned a higher service level. For example, one rule might be to set the service level at 90 percent (i.e., risk of stock-out of 0.10) for items with unit prices greater than \$1,000, and at 99 percent for all other items. An item is then stocked if the bootstrap methodology for a given risk assigns a nonzero ROP.¹⁶ Chapter Four will explore this idea further.

¹⁶This is the simplest stock rule. One could augment this rule for inexpensive, low-demand items that are assigned a zero ROP by allowing a nonzero RO.

REDUCING THE LENGTH AND VARIABILITY OF THE REPLENISHMENT PROCESS

The greatly reduced inventories shown in this chapter, and the greatly increased service to the customer to be shown in the next chapter, are a function of two improvements:

- Application of the bootstrap and dollar-banding techniques to “rightsize” the inventory.
- Assumed continued improvement in replenishment order and ship processes so that they are less variable; in particular, for today’s worst 10 percent, which currently take longer than 30 days, the process is improved so that all (non-back-ordered) replenishments take 30 days or less.

The Marine Corps, as part of the Precision Logistics initiative, aggressively pursues improvements in the order and ship process for wholesale supplies. This is certainly important for the many maintenance customers awaiting repair parts from wholesale sources, as we showed in Chapter One in our discussion of the “wholesale order and ship time penalty” imposed on repair times when an item was not available from the SMU. But, with the adoption of the ideas presented in this chapter, we firmly expect the dependence of most Marine customers on *direct* shipments from the wholesale system to decrease greatly.

The SMU will continue to be the buffer between the wholesale system and Marine customers, and the efficiency and rapidity of the wholesale order and ship process for SMU replenishments will, in fact, become increasingly vital. The faster replenishments come, the broader the stocks the SMU can support and the greater its reliability in having items on the shelf when the customer comes demanding. To do so, the SMU must take steps to make sure this replenishment process is both *faster and less variable*.

To understand this, recall how the bootstrap works. It estimates for a given lead-time distribution (i.e., replenishment time) the probability that no more than some number of customer demands will arrive at the SMU; that is, it adds stock to cover the risk of not having enough if demands turn out to be unusually high while it is awaiting its replenishment order to arrive from the wholesale system. How-

ever, just as the number of demands can vary within that lead time, so too can the lead time itself vary. We used lead times that were capped at 30 days. But obviously, not all replenishment orders arrive within 30 days. In fact, for I MEF replenishments, only about 90 percent of them tend to arrive within 30 days.¹⁷ This means that 10 percent of the time the lead time (OST) for I MEF replenishments will exceed 30 days and the SMU is at greater risk of stocking-out when customer demands arrive than if the OST was 30 days or fewer. Thus, that one in 100 risk (when buying to $r = 0.01$) is actually greater than one in 100.

This is because of the variability of the replenishment order and ship process. Figures 3.3 and 3.4 illustrate this for I and II MEF SMU replenishment requisitions. They show the distribution of OSTs for

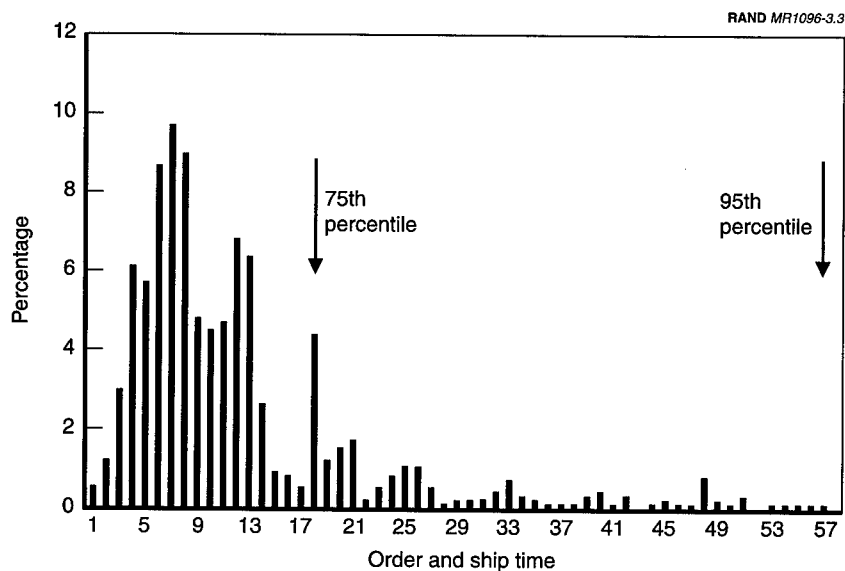


Figure 3.3—Distribution of Replenishment OSTs for I MEF SMU (MMC100), CY 1998

¹⁷This is for immediate issues from DLA depots only. Back-orders and direct vendor deliveries are excluded.

replenishments filled out of DLA depots (with no back-orders) experienced in calendar year 1998. (For the sake of readability, and because there is an extremely long right-hand “tail” to the distribution, we do not show the worst five percent of the cases; i.e., we cut the right-hand side of the chart at the 95th percentile.) The 75th percentile for I MEF requisitions is 18 days; for II MEF (see Figure 3.4) that figure is 15 days. But many requisitions take more than that—25 percent in fact. We could increase our guarantee against a stock-out by buying further to the right of the distribution.

But each move to the right increases the depth of the stock required and therefore the inventory investment in dollars. If, say, we wanted 95 percent confidence about covering the order and ship lead time, we would set it at 57 days (for I MEF) and 37 days (for II MEF). That would lead to hugely increased ROPs, to ensure that we would cover the vast majority of demands that *might* come in during those 57 or 37 days. And greater depth of stock means less breadth, which

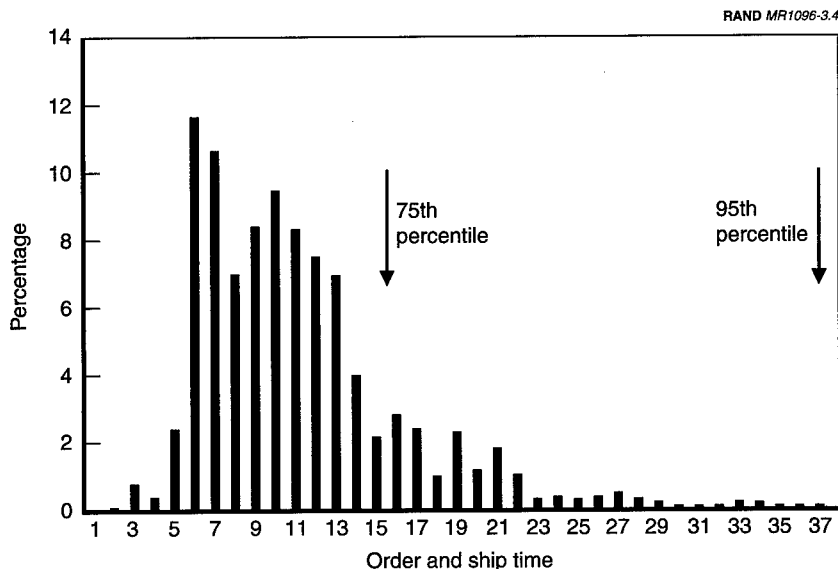


Figure 3.4—Distribution of Replenishment OSTs for II MEF SMU (MML100), CY 1998

means smaller fill rates and ERO fill rates and worse performance for the customer.¹⁸

Calculating the cost of the inventory when current lead-time performance is included in the bootstrap risk calculation drives home the need to reduce OST. Table 3.4 shows the inventory cost when a maximum 30-day OST is assumed and when the actual OST distribution is used. The result is that it costs \$20 million more to achieve the same risk of stock-out with today's long and variable OSTs.

The answer is not to buy more security by increasing the depth of the stock; it is to gain more security by reducing the unpredictability of the process and shortening the overall time it takes. That is, instead of buying to a 95th percentile of 57 or 37 days, the SMUs should strive to make the 95th percentile *18 days*—or even less.¹⁹ By so

Table 3.4
Comparison of Bootstrap Inventories with Different
Lead-Time Distributions

	Bootstrap GABFs, $r = 0.01$, OL = 30 DOS	
	Maximum 30-Day Lead Time	Actual, Variable Lead Time
Number of lines	32,645	32,645
Value at the ROP	\$13.2	\$34.7
Value at the RO	\$16.9	\$38.1

NOTE: Dollar figures are expressed in millions.

¹⁸These distributions are far worse if back-orders are included. We argue against doing so, and against basing inventory levels on OSTs including back-orders, partly because of the vast expense (or reduced customer service levels) that would result and also because it is virtually impossible to predict on a systematic basis which items will be stocked-out in the future at the wholesale level. Therefore, we argue that the SMU should accept the risk of a stock-out at the wholesale level and so accept that in those unfortunate cases it simply will not have access to the stocks it will need. Certainly, if there are a few items for which a stock-out at the retail level is simply not acceptable, then the SMU can greatly increase its safety stock to cover the expected period of empty shelves at the wholesale level.

¹⁹As a result of the implementation of Velocity Management (VM), the Army's equivalent of Precision Logistics, replenishment times in such units as the 101st Airborne (Air Assault) Division and the 82nd Airborne Division consistently produce 95th percentiles substantially below 18 days for non-back-ordered repair parts. This reduction—and the introduction of dollar banding—has allowed the 101st to reduce its inventory investment by more than half while increasing service to its customers.

doing, customer satisfaction will increase (through fewer stock-outs) at no greater cost.

As an example, see Figure 3.5, which shows the distribution of replenishment OSTs for the III MEF SMU on Okinawa, Japan. This figure looks rather different from the previous two; in fact, it is distinctly bimodal. In the first half of 1998, the Marine Corps took steps to improve the traditionally very poor order and distribution times for III MEF requisitions. Much progress was made internally, improving requisition issuing and receipt processing steps. A major change, however, was to work with DLA to set up “depot direct” delivery service, making more use of premium air transportation. As a result, the distribution of OSTs shifted, as the figure shows. The median dropped substantially, from 42 to 17 days, and the 75th percentile also showed strong improvement, falling from 53 to 40 days.

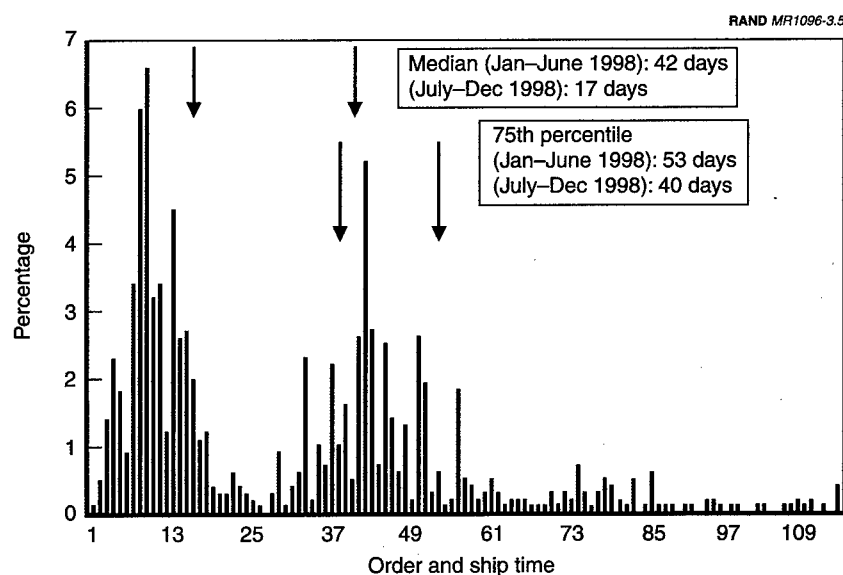


Figure 3.5—Distribution of Replenishment OSTs for III MEF SMU (MMR100), CY 1998 (Change in Median and 75th Percentiles Between First and Second Half of Year Noted)

As this example suggests, order and ship process improvement is not completely in the hands of the SMU. But much of it is. As the quarterly Marine Corps wholesale OST reports demonstrate, much of the length and variability of the SMU replenishment process lie in on-base activities—issuing requisitions and posting receipts to inventory. Were the SMU to apply Precision Logistics methods to these parts of the process and continue to work with DLA and the wholesale system on the other parts of the process, it would likely see a substantial improvement in the overall length and predictability of the replenishment process. This can only mean better service to its many customers. This kind of improvement, integrated into application of the bootstrap determination of ROPs, can have substantial impact on the III MEF SMU's ability to support its customers and raise its low ERO fill rates with no additional inventory investment—indeed, with less.

CONCLUSIONS

This chapter presented two new techniques to improve the Marine Corps retail inventory performance, bootstrapping the ROP and dollar banding the inventory. Taken together, these two techniques prescribe new ways for the Marine Corps to set inventory levels. We have also argued that reducing the OSTs, and the variability of the OST distribution, will allow the SMU to provide better service at a reduced inventory cost. In the remainder of this report, we discuss and apply a new methodology, the Virtual SMU, that will help inventory managers evaluate exactly how these new techniques will perform in practice.

Some Marine Corps leaders may object to the dollar-banding concept on the basis that the inventory should be set according to equipment warfighting importance rather than cost. This is a significant and real concern that we have not overlooked. In our analyses, we had item cost data available, but we did not have an accurate measure of items' warfighting importance.²⁰ The advantage of the bootstrap ROP methodology, however, is that service levels can be specified according to any criteria; we used cost, but the inventory

²⁰We did have CECs available, but these are widely acknowledged to poorly represent the warfighting importance of items.

manager might just as well use some measure of warfighting importance if it is available. For example, the bootstrap ROP can be readily adapted to support the ILC Initiative's "Quadrant Model." As the next chapter will show, we successfully used cost to dramatically improve readiness, as measured by the ERO fill rate, which argues that there is merit to this approach, at least compared with the current methodology. However, it is reasonable to expect further performance improvement using a combination of cost and warfighting, repair, or equipment importance, appropriately targeted toward the desired dimension of improvement.

PUSHING THE LEVER: DEMONSTRATIONS OF IMPROVED INVENTORY PERFORMANCE

Chapter Two highlighted some of the shortcomings of the current Marine Corps methodology. Chapter Three demonstrated the bootstrap and dollar-banding techniques abstractly but provided little indication of how a new inventory based on these techniques would actually perform under realistic conditions. The Virtual SMU, which we briefly discuss below (and in more detail in Appendix E), was thus created to evaluate this question. Using the Virtual SMU, we demonstrate that our new techniques, in conjunction with ongoing Precision Logistics efforts, would likely show significant improvement in inventory performance. In particular our results suggest inventories that promise to

- perform as well as the existing inventory in terms of fill rate, satisfaction rate, accommodation rate, and ERO fill rate—at less than one-half the current inventory cost, and
- outperform the existing inventory, achieving ERO fill rates of better than 70 percent and fill rate metrics near or better than 90 percent at three-fourths of the current inventory value at the RO.

THE VIRTUAL SMU

We developed the Virtual SMU to test various policy alternatives “off line” but under the rigors of actual customer demand patterns experienced by I MEF. Its purpose is to evaluate how new inventory management techniques could perform in practice, with the goal of selecting inventory management policies, techniques, and algo-

rithms that perform well. We use it to answer the question, "How can the SMU achieve the best possible *overall* performance (as measured by such inventory metrics as fill rate and ERO fill rate) within imposed funding constraints?"

As discussed in detail in Appendix E, the Virtual SMU performs all the standard accounting procedures that any inventory management system should. It starts with an initial inventory balance, subtracts demands and/or accounts for due-outs as requisitions are presented, orders and adds replenishments as needed, and otherwise tracks financial constraints and calculates inventory performance metrics. While simple in concept, the code is fairly complicated because it allows for a variety of policy alternatives. Furthermore, it is "open" in the sense that it is uncompiled, which allows further policy options to be entered as modifications to the code.

The key idea of the Virtual SMU is that it "plays back" the actual demands experienced by the SMU when evaluating policy alternatives. The idea is that, with a historical archive of data of, say, two years, one can analytically turn the clock back to a year ago. The first year's data are used to establish inventory policy parameters, from which a "virtual" simulated inventory is developed, and this inventory is then tested using the second year's data. Thus, the performance of a particular inventory policy can be evaluated under conditions as close to actual as possible.

CREATING A NEW GABF USING THE BOOTSTRAP

We applied the bootstrap methodology to demand data for February 1997 through January 1998 from the I MEF to generate alternate ROs and ROPs that reflect an improved inventory posture of how the I MEF General Account could be stocked. For the sake of simplicity, in this section, the RO was set using an operating level based on 30 days of supply and the on-hand quantity was set randomly between the RO and ROP; the details are discussed in Appendixes D and E.

Having defined new ROs and bootstrap ROPs, a new GABF was created and evaluated using the Virtual SMU. This required additional assumptions for transaction information that was not available, such as the specific order that requisitions arrived at the General Account and the simulation of the stock replenishment process. To account

for this, we simply assumed that each day's high-priority demands arrived first and were filled first. Then, as in Appendix E, we "played back" the actual requisitions presented to the SMU for the period May 1 to December 31, 1998.

We gave the SMU a \$7 million budget for each quarter, which is roughly the actual quarterly budget. We set aside a varying percentage of the total budget to cover pass-throughs,¹ depending on the number of lines and the risks set for the ROPs. We specified that replenishment buys would be placed every seven days and allocated one-thirteenth of the available quarterly funds to be spent each week. We also used a simple buy rule:² Weekly, rank all NSNs with inventory positions less than their respective ROPs, from least expensive to most expensive. Starting with the least expensive, buy the quantity necessary to return the inventory position to the RO, proceeding down through the ranked NSNs until the week's allocated funds are expended.

Table 4.1 shows that using the bootstrap GABF with $r = 0.01$ increased the fill rate from 72 percent (using the original GABF) to 85 percent and the ERO fill rate³ increased from 54 percent to 68 percent. These results were achieved with an on-hand inventory value significantly less than the existing inventory value, and roughly equivalent investment costs at the RO and ROP.⁴ Perhaps the most

¹A pass-through is a requisition for an item not stocked locally. The requisition is literally passed through to the wholesale system.

²See Appendix E for a discussion of this buy rule and an evaluation of its performance. Briefly, we found that this buy rule, without any other changes in SMU inventory policy, improved inventory fill rates while removing the burden of making manual buy decisions.

³Note that the ERO fill rate calculated in these experiments differs from the actual ERO fill rate because secondary repairables stocked at the RIP are not included in the demand stream, only consumables. In spite of this, the performance of the actual GABF in these experiments is expected to be very similar to the performance of the inventory in actual practice.

⁴Of course, as was originally discussed in Chapter Three, the demand stream we used was not the complete demand stream used to calculate the actual GABF. If it had been, then the value of the bootstrap GABF at the ROP would have been somewhat larger. In spite of this, the actual GABF, sized to handle a *larger* demand stream, still performed worse than the smaller bootstrap GABFs, with the same demand stream used to measure both types of GABFs.

Table 4.1
Comparison of the Actual GABF Performance with
Various Bootstrap GABFs

	Actual GABF	Bootstrap GABFs			
		$r = 0.05$ OL = 30 DOS	$r = 0.03$ OL = 30 DOS	$r = 0.01$ OL = 30 DOS	$r = 0.01$ OL = 90 DOS
ERO fill rate	54%	52%	58%	68%	66%
Fill rate	72%	73%	78%	85%	85%
Satisfaction rate	81%	87%	90%	92%	92%
Accommodation rate	83%	84%	86%	92%	92%
Number of lines (NSNs) stocked	13,159	14,014	17,682	32,645	32,645
Value at the ROP	\$13.5	\$4.7	\$6.7	\$13.2	\$13.2
Value on-hand	\$25.5	\$5.1	\$7.2	\$13.7	\$21.3
Value at the RO	\$23.6	\$6.4	\$8.8	\$16.9	\$24.5

NOTE: The number of lines and inventory value figures are repeated from Table 3.1. Dollar figures are expressed in millions.

striking contrast comes with the comparison of the $r = 0.05$ bootstrap GABF with the actual GABF. This bootstrap GABF showed very similar performance to the actual GABF in all of the metrics in Table 4.1 using an inventory with a fraction of the dollar value of the actual GABF. *Essentially, this bootstrap GABF achieved equivalent performance for less than a third the price of the actual inventory.*

Table 4.1 also shows the performance improvements that can be achieved by the bootstrap GABFs for various risk levels, where we see that as the risk is reduced, the various inventory metrics increase accordingly. For example, the $r = 0.01$ bootstrap GABFs have an accommodation rate of almost 92 percent, which means that 92 percent of all items requested were supposed to be stocked (in contrast to 83 percent in the actual GABF). Similarly, 92 percent of the demands that were supposed to be stocked were actually filled out of local stock (the satisfaction rate) with the $r = 0.01$ bootstrap GABFs, while only 81 percent were with the actual GABF. In particular, the GABF using $r = 0.01$ with 30 DOS OL outperforms the existing inventory, achieving ERO fill rates of almost 70 percent and fill rate metrics near 90 percent for less than three-fourths of the current inventory value at the RO.

FURTHER IMPROVEMENTS WITH DOLLAR BANDING

Inventory performance can be further improved by targeting various classes of parts, perhaps by either dollar value (expensive) or some measure of combat or functional importance (nonessential), and appropriately allowing an increased risk level for items in these groups. That is, using a risk level of 1 percent for all parts is to treat everything as highly important when, in fact, groups of parts could certainly be allowed a higher risk with little detriment to the inventory's performance.

Table 4.2 compares some dollar-banded alternatives with the non-dollar banded bootstrap GABF with $r = 0.01$ and 30 DOS OL from Table 4.1. The non-dollar banded bootstrap GABF is listed in the second column. The next column over shows that significant savings can be achieved with little degradation in performance by increasing the risk on expensive items. In this particular example, the risk was increased from $r = 0.01$ to $r = 0.5$ for items costing more than \$1,000. For all other items the risk was maintained at $r = 0.01$. This modification saved \$2.1 million in inventory investment at the ROP, and \$3 million at the RO, while maintaining the ERO fill rate and decreasing the overall fill rate by only 4 percent, respectively.

The three right columns in Table 4.2 then show how additional increases in the ROP for inexpensive items can improve ERO fill rate and fill rate performance while keeping the inventory investment equivalent to the non-dollar banded alternative. Here, we set very high service levels for inexpensive items as a cheap hedge against unusually large, unforeseen future demands. Note that the safety level (via the ROP) is increased and not the operating level. For these three, the bootstrap ROP risk is set at $r = 0.001$ for all NSNs below a threshold, starting at \$1 in the middle column, then \$5, and then \$8 on the far right. The alternative in the last column approaches the inventory cost of the non-dollar banded GABF but shows appreciable improvement in inventory performance. In particular, both the accommodation and satisfaction rates are at 93 percent, giving a fill rate of 87 percent and an ERO fill rate of just over 75 percent.

The last alternative begins to shed some light on the limits of what is achievable, as well. That is, the accommodation rate is strictly a

Table 4.2
Comparison of the Performance of Various Bootstrap GABFs

	<i>r</i> = 0.01 (no dollar banding)	Dollar-Banded, Bootstrap GABFs			
		<i>r</i> = 0.01 for UP < \$1K; <i>r</i> = 0.5 for UP ≥ \$1K	<i>r</i> = 0.001 for UP < \$1; <i>r</i> = 0.01 for \$1 ≤ UP < \$1K; <i>r</i> = 0.5 for UP ≥ \$1K	<i>r</i> = 0.001 for UP < \$5; <i>r</i> = 0.01 for \$5 ≤ UP < \$1K; <i>r</i> = 0.5 for UP ≥ \$1K	<i>r</i> = 0.001 for UP < \$8; <i>r</i> = 0.01 for \$8 ≤ UP < \$1K; <i>r</i> = 0.5 for UP ≥ \$1K
ERO fill rate	68%	68%	71%	73%	75%
Fill rate	85%	81%	82%	85%	87%
Satisfaction rate	92%	86%	88%	92%	93%
Accommo- dation rate	92%	91%	91%	92%	93%
Number of lines	32,645	32,537	32,537	32,537	32,537
Value at the ROP	\$13.2M	\$11.1M	\$11.4M	\$13.0M	\$14.1M
Value at the RO	\$16.9M	\$13.9M	\$14.2M	\$15.8M	\$17.0M

NOTE: Comparison of the bootstrap GABF with $r = 0.01$ and OL = 30 DOS from Table 4.1 to various dollar-banded bootstrap GABFs. UP = unit price.

function of stocking the right items—those that are requested—for which the selection of about 33,000 items out of approximately 85,000 seems to be an upper bound. Said another way, to achieve the last 7 percent of the accommodation rate would require stocking the remaining roughly 52,000 items, which is not likely to be fiscally possible or prudent. Similarly, improvements on the 93 percent satisfaction rate are likely to have a steep curve of diminishing returns. Since the fill rate follows immediately from the accommodation and satisfaction rates, for all practical purposes, it is similarly bounded at about the 85–90 percent range.

CONCLUSIONS

This chapter has demonstrated that significant improvement in inventory performance is achievable using the bootstrap ROP and dollar-banding methodologies. In particular, we have shown that dramatic increases in the ERO fill rate can be achieved through

improvements in supply policies and procedures and with a reduction in wholesale OSTs. We have also shown that improvement in the ERO fill rate requires broadening the inventory. In this chapter, we have almost tripled the number of types of items stocked in the inventory, from almost 14,000 to around 35,000.⁵ Under budgetary constraints, such broadening entails reducing the quantity of the items currently stocked. To do this, the Marine Corps *must* revise its current methodology to incorporate more precise measures of demand variability in its inventory calculations. The bootstrap ROP methodology is one technique for doing this.

A distinct advantage of the bootstrap ROP and dollar-banding techniques we have presented is that the problem of defining what and how much to stock is reduced to terms that the inventory manager naturally thinks in and that the inventory manager can make rational judgments about. Specifically, an inventory manager's job is to stock parts to achieve a desired service level—that is a problem they worry about day-in and day-out: “Will I have enough stock to meet demand?” Using the bootstrap ROP methodology, the inventory manager is freed from arbitrary rules, such as days of supply, and can specify by item the service level desired. Under dollar banding, we used unit price as the criteria for setting service levels, though the methodology can just as easily be applied using the ILC Initiative's “Quadrant Model” to specify service levels.⁶

⁵It is, of course, an open question whether the supply infrastructure, such as warehouse size and manning, can support such a broadening of stock. Implementation of this approach would require some study of this aspect of the problem. However, our initial expectation is that appropriate planning and implementation would allow significant broadening with minimal impact on operations. After all, the warehouse operations will still handle roughly the same number of items; it is only a question of whether the items are given a bin for local storage. That is, the SMU must handle all of the approximately 30,000 items that experience some annual demand. It is mainly a question of whether 15,000 items are stocked in bins and the other 15,000 are processed and handled as pass-throughs, or whether all 30,000 are given bins.

⁶For example, items that are stocked by the Marine Corps and classified as “bottleneck” or “critical” would have a high service level assigned; items classified “routine” or “leveraged” would have a lower service level assigned. Such a scheme might be further refined using dollar banding so that within a particular category, “critical” for example, less expensive items are given a higher service level than more expensive items.

Of course, the adoption of simple cost-based rules in a military system may not always be appropriate. The major point, however, is that risk of stock-out may be assigned to classes of parts using the bootstrap methodology. Price classification is only one such approach; military importance might also be used. Additionally, the current use of a fixed 60 DOS for the operating level seems excessive. ROPs with appropriately sized safety levels, as with the bootstrap ROPs, should allow the inventory to function with a smaller operating level yet with a higher confidence of fewer stock-outs. Smaller operating levels require smaller incremental outlays of funding to buy items from the ROP back to the RO, and result in ancillary benefits, such as smoothing work flow in the warehouse.

**FUTURE STEPS: IMPLEMENTATION,
METHODOLOGICAL IMPROVEMENTS,
GENERALIZATIONS**

Our methodology for improving Marine Corps inventory performance is based on modern computational methods. The bootstrap ROP technique represents an advance in both actual Marine Corps practice and in basic inventory theory. Our techniques, while illustrated using I MEF General Account data, are broadly applicable across the Marine Corps, including the RIP (repairable issue point) and the ICP (inventory control point).

Bootstrap calculations, while computer intensive, are transparent to the inventory manager. Furthermore, as we previously mentioned, the advantage of the bootstrap ROP over current Marine Corps practice is its input in terms of desired service level—a natural way for an inventory manager to specify inventory requirements.

We have shown through simulation with actual data that the resulting inventory performs significantly better. This result was demonstrated with 1st Marine Expeditionary Force data and would likely hold for other Marine Corps organizations. Indeed, most organizations for which the standard assumptions of inventory theory do not work well will benefit from application of the bootstrap ROP methodology. Below, we discuss some specific implementation ideas and requirements for the Marine Corps and for other organizations.

FOR THE MARINE CORPS

This report has focused on unilateral measures that the Marine Corps supply organization can take to improve inventory performance in general and the ERO fill rate in particular. The ERO fill rate measures inventory effectiveness from the perspective of the mechanic. It has an intuitively direct connection to unit readiness. We have shown that ERO fill rates of around 75 percent are achievable following the adoption of new inventory management techniques, such as the bootstrap ROP and continued improvement in OST performance. Further improvement beyond this, however, will require cross-functional cooperation between the SMU, the maintenance providers, and the using units. This is because the ERO fill rate is not only a function of supply and SMU inventory but also a function of how parts are ordered and repairs inducted; communication, cooperation, and innovation in the true spirit of process improvement will ultimately be required.

Implementation. The concepts and tools described in this report are available for "customized" implementation in the Fleet Marine Force now. Two main steps are required for successful assimilation of these ideas and techniques into the Marine Corps: (1) test implementation demonstrating proof of concept in actual practice, followed by (2) formal incorporation into USMC systems and processes. In step one, we believe that the concepts should be subject to additional refinement in the Virtual SMU with Marine Corps input. This should be followed by a real-world test at an SMU where the new inventory guidelines are implemented, their performance observed, and adjustments made as necessary to accommodate SMU needs and requirements. Such a test could then proceed to a rolling, customized installation at each MEF. In step two, given a successful demonstration or rolling implementation, the Marine Corps would then formally promulgate instruction changes, incorporate the techniques into standard computer hardware and software,¹ and update training curriculums. As mentioned previously, this implementation will be most successful when coupled with an aggressive

¹The current RAND software, while sufficient to support the test implementation, is insufficient and inappropriate for long-term Marine Corps field use. Under formal implementation, these concepts, algorithms, and techniques should properly become an integral part of standard USMC systems, perhaps such as ATLASS.

SMU effort to reduce OSTs. We suggest that implementation of the bootstrap ROP methodology is a natural endeavor to incorporate in the ILC Initiative.

Methodological Improvements. The implementation and application of the bootstrap ROP methodology in this report was the simplest possible. We did not try to incorporate additional knowledge that Marine Corps logisticians and users have about specific item usage and other general usage patterns. Thus, many enhancements to the method could be considered in an actual implementation. For example, the method could be modified to account for seasonal items and items ordered as a group. As was discussed in the text, we also did not attempt to set service levels by anything other than cost (through dollar banding), but additional measures could be developed and incorporated to account for the "importance" of an item. We also used a simple DOS approach for setting operating levels. More sophisticated techniques, such as economic order quantity (EOQ), could be employed to balance the costs of ordering versus holding stock.

Generalizations. In addition to improving General Account and other consumable inventory management practices, we see great promise in extending these ideas to other inventory problems. For example:

- The bootstrap methodology can be generalized to more complicated inventory problems, such as stocking repairable items at the RIP or the ICP. For repairables, an issued item is coupled with a returned broken item that may be repaired or refurbished and returned to stock. In spite of this more complicated scenario, a bootstrap scheme can be readily implemented that would account for these additional requirements and still produce a simple risk distribution from which ROPs could be set in the manner we have just presented.
- Furthermore, some of the dollar-banding ideas could be applied to building parts blocks for deploying Marine Expeditionary Units or for Maritime Prepositioning Ships. The concept can also be generalized to account for weight or volume when financial constraints are not an issue. For example, parts blocks usually favor stocking parts with smaller volume over parts with

greater volume. Bootstrapping techniques could also be applied to deployment parts demand histories to generate suggested stock levels.

FOR OTHER ORGANIZATIONS

The metrics and inventory management ideas that we have discussed, though Marine Corps specific, are easily generalized to other organizations. Like the other military services, the Marine Corps faces an unusually difficult task—having to supply a large and diverse array of parts for very specialized equipment. Other organizations with less-demanding inventory requirements (in terms of inventory size, budgetary constraints, smoother or more reliable demands, etc.) can expect similar and, most likely, better results.

In addition to the bootstrap results, organizations that incur significant costs (either monetary, time, or some other type of cost) when the local inventory fails to stock a part could benefit from the application of a metric resembling the ERO fill rate. Such an all-or-nothing metric only gives “credit” when all parts are supplied. The value of this type of metric is its change of the focus from evaluating inventory performance at the individual part level to an aggregate level that is of organizational importance.

Implementation. The basic details of the bootstrap ROP methodology can be found in Appendix D. Its application requires determining the specific implementation details, such as the period to resample, and the availability of demand and lead-time data. This in turn requires the evaluation of autocorrelation in the demand stream and, perhaps, correlation among items. Thus, while the basic calculations are relatively straightforward, the proper application of the methodology and the evaluation of its performance are nontrivial and should only be carefully undertaken.

Methodological Improvements. As with the Marine Corps, additional knowledge that the organization’s logisticians and users have about specific item usage and other general usage patterns (such as seasonal items and items ordered as a group) could be incorporated. For the commercial sector, setting service levels by cost through dollar banding might be particularly appropriate, though additional measures could be used to account for the “importance” of an item.

Generalizations. As was discussed with the Marine Corps, the bootstrap methodology is very flexible and can be extended to many nonstandard and complicated inventory problems. Some of these are discussed at the end of Appendix D.

CRITICAL REPAIRS AND THE ERO FILL RATE

LIMITING THE ERO FILL RATE TO CRITICAL REPAIRS

Part of the reason for the disconnect between the supply and maintenance functions is the supply system's rather opaque view of the maintenance system's most important needs. The only information the supply system typically gets from the mechanic is the priority of a requisition or the combat essentiality code (CEC) of the part needed. Neither of those, however, gives the supply system the fullest visibility of how well it is helping mechanics achieve their main goal: returning critical systems to operational status in the shortest time possible.

To make that connection better, the Marine Corps has adopted the "ERO fill rate" metric, which measures the percentage of critical repairs that receive *all* their high-priority parts from local supply with no back-orders.

The ERO fill rate is limited to critical repairs. Excluded from the measure are maintenance actions with no time criticality or whose need for parts is not intense. Identifying what is or is not a "critical repair" is not straightforward, however.

The standard Marine Corps Integrated Maintenance Management System (MIMMS) includes the "maintenance category type" (typically referred to as "catcode"), which identifies the criticality of the repair action. Repairs on deadlined weapon systems are marked

with catcodes of M (for MCGERR-reportable¹ equipment) or P (for deadlined equipment which is not MCGERR-reportable).

Unfortunately, in the MIMMS archive the maintenance category code (hereafter "catcode") records the *last* catcode assigned to an ERO. That is, for example, an ERO might open in N status (nondeadlined), stay that way for 100 days, go into M status on day 101 (say if a deadlining fault was found), and close out the same day with the final catcode (captured in the archive) being M. Clearly, we would not want to count this 101-day repair action as MCGERR-reportable deadlined, because only one day counted as deadline time. Alternatively, an ERO could have catcode M for 50 days and N for the final day and be counted as an N. By simply looking at catcodes, we might miss this repair.

Instead, we rely on information contained in the fields "catmdays" (days deadlined for MCGERR-reportable systems) and deadline date (the date a system was deadlined). As it turns out, both are necessary for selecting the right cases.

Take MCGERR-reportable systems first.

Catmdays is a convenient measure, if used properly, of the time during a repair action that a system is deadlined. It is possible, though, that the repair action time may be greater than the deadlined time. We face two choices: only count catmdays in all repair actions or only include repair actions that meet some threshold of catmdays. We choose the latter.

Consider two cases:

- Repair One has two EROs from two different units with a repair action time of 100 days and catmdays of 50 days.
- Repair Two has two EROs from two different units with a repair action time of 100 days and catmdays of 98 days.

¹MCGERR stands for Marine Corps Ground Equipment Resource Reporting. Specific types of equipment, identified by TAMCN (Table of Authorized Materiel Control Number) by policy have their readiness reported on a weekly basis. Readiness reporting policy is set out in Marine Corps Bulletin 3000.

In the first case, while we would like to report on the repair time when the system was deadlined, we don't know to when to assign it; after all, for 50 days the system was not deadlined. In the second case, repair action time and catmdays are virtually identical so we can safely assign the whole repair as a "critical" repair of a deadlined system.

Thus, our first selection filter is:

Select all repair actions where catmdays is at least 90 percent of the repair action time.

When we looked closely at the data, however, we found many cases in which the final catcode was M but catmdays was zero; however, in almost all of those cases the deadline date field was filled in. We use this as an alternative means of selecting cases. We simply use deadline date as a surrogate for catmdays:

Select all repair actions in which the time between deadline date and the final close date is at least 90 percent of the repair action time and the catcode is M.

We put in the additional criterion of an M catcode because that increases the probability that the entire repair action was on a deadlined system. That is, the system was deadlined close to the initial date received in shop (DRIS) and closed out as an M (deadlined) repair status. If the catcode were different (e.g., N), that would not be usable because, although the deadline date was close to the DRIS, we would have no idea when the system came out of deadline status—it might have been one day after the DRIS. So while this assumption is not perfect—it is possible to be deadlined early, come out of deadline status, and go back into it just before the end of the repair action—we think it does a fairly good job of capturing most or all important repairs and limiting the inclusion of noncritical repairs.

Finally, repair of deadlined non-MCGERR-reportable systems.

This logic follows the previous one very closely. Non-MCGERR systems do not have catmdays reported but they do have deadline dates. Therefore, for non-MCGERR deadline repairs we use the following logic:

Select all repair actions in which the time between deadline date and the final close date is at least 90 percent of the repair action time and the catcode is P.

We put in the criterion of a P catcode for the same reason we require an M catcode above.

DEFINING THE ERO FILL RATE

The ERO fill rate, which is the fraction of EROs for important PEIs and secondary repairables that have all of their critical parts immediately available from the SMU, is a new inventory metric. It measures supply performance in supporting the repair of equipment that affects combat readiness and is essential to the mission of the Marine Corps. This section explains how the ERO fill rate metric is computed and discusses why it is important.

ERO fill rate is the percentage of critical EROs for which all high-priority requisitions were immediately available from local supply (i.e., the General Account, the RIP, and their forward locations).

A critical ERO is one involved in the repair of a deadlined weapon system (MCGERR and non-MCGERR-reportable) or a secondary repairable (with catcode D, F, H in echelon of maintenance 3 or 4). An ERO is part of a deadlined weapon system repair if the deadline time is at least 90 percent of the total RCT, from first DRIS to last close-date, measured using either catmdays or the difference between close date and deadline date. These are the same EROs reported in the RCT metric. A high-priority requisition has an 06 or better priority code.

The motivation behind the ERO fill rate metric is that a repair mechanic, after determining the parts required to fix a deadlined PEI or broken secondary repairable, should be able to get all critical repair parts out of retail inventory immediately and effect immediate repair of the deadlined equipment. The goal of the metric is to capture supply performance in facilitating the repair of equipment that affects combat readiness.

The metric differs from other inventory measures in that it only gives "credit" for supplying all the parts for a repair. If one critical part is unavailable, it is unlikely that the repair can be completed. So, get-

ting 90 or 95 percent of the parts is not sufficient. Only 100 percent of the parts will complete the repair, and that is what ERO fill rate measures.

Calculating the ERO Fill Rate

Calculating the ERO fill rate requires merging MIMMS data with SASSY data for critical EROs and high-priority requisitions. The process is conducted in the following general steps:

From MIMMS select the EROs that closed in a quarter that were part of critical repairs. Only closed EROs are used to avoid any bias that would result from including EROs that have not received all of their parts and to prevent double counting open EROs from time period to time period. Closed EROs are defined to be those with status codes 03 and 15.

From SASSY select all requisitions with priority greater than or equal to 06. For each requisition determine whether it was provided immediately from local supply by examining the "AE1" records. If all AE1 records list a BA status and at least one shows an RIC for an SMU, the requisition is classified as immediate retail issue, otherwise it is not. The priority restriction screens out less-critical requisitions and those that may have been attached to the EROs out of convenience.

Merge the ERO records with the SASSY records by ERO number, MEF, and whether the requisition document date falls between the ERO DRIS date and the ERO close date. Delete canceled requisitions and all EROs that had no requisitions so that the population under consideration finally consists of important EROs and their associated high-priority requisitions for a given period.

Finally, evaluate each ERO's set of requisitions to see if all requisitions had a BA status. The ERO fill rate is the proportion of filled EROs out of the whole population of EROs that had at least one requisition that matched between the MIMMS and SASSY databases.

BASIC INVENTORY THEORY AND TERMINOLOGY

Before we can discuss improvements to supply policy, it is necessary to introduce some basic inventory theory concepts and standardized terminology. While this appendix is a bit dense and abstract, clearly defining concepts and terminology is important for establishing a common basis of reference and understanding.

TRACKING STOCK LEVELS

Standard inventory theory describes establishing a local inventory in terms of determining which items to stock and then specifying the amount of each item to stock. See Tersine (1994) or Arrow, Karlin, and Scarf (1958), for example. For each stocked item the following three quantities are tracked over time:

- On-hand stock (OH), which is the actual amount of the item in the warehouse.
- Due-in stock (DI), which is the total amount of stock that has been ordered from the next-higher echelon of supply to replenish the on-hand stock plus items due in from repair (as applicable).
- Due-out stock (DO), which is the total amount of unfilled requests resulting, for example, when the local inventory runs out of the item. Due-outs are also referred to as back-orders by customers.

From these three quantities the inventory position (IP) is calculated. For a given item the IP is

$$IP = OH + DI - DO.$$

That is, the inventory position equals the on-hand stock plus the due-in stock minus the due-out stock. In essence, it is the instantaneous balance if all accounts (due-in and due-out) were immediately settled.

A stocked item is characterized in terms of two quantities: the requisitioning objective (RO), which is the maximum on-hand quantity, and the reorder point (ROP), which is the point at which new replenishment stocks are ordered. The ROP is always less than the RO for stocked items. The IP is used with the ROP and the RO to determine when and how much to order. In a system with automatic reordering, when the IP hits or drops below the ROP an order is placed for the quantity necessary to bring the IP back up to the RO. Thus, automatic reordering maintains the IP in between the RO and the ROP at all times. The inventory literature refers to this as an (s, S) system.

For example, one possibility is to set the ROP at one unit less than the RO (an “[s,s-1]” policy in inventory theory parlance). In this case, with the inventory starting out at the RO, a replenishment order would be generated immediately after any demand is received at the SMU because, even if the demand were for one item, the IP would hit the ROP and immediately trigger a replenishment buy order. In this case, the IP always equals the RO, sometimes called the “stocked level.”

Figure B.1 illustrates the relationship between on-hand stock and IP over time for a hypothetical item. Initially, the item starts out with the IP at the RO but with the on-hand stock (in gray) at the ROP—a replenishment order has just been placed. Over time the two decrease as demands are received and, before the next replenishment order is placed, the item goes into back-order status, with the on-hand stock reaching zero and the due-out quantity increasing. Not long after that, however, a replenishment order comes in, pushing the on-hand stock back into the positive, and the IP and on-hand quantities coincide. A short while later, the IP reaches the ROP and a replenishment order is placed, and the IP jumps back up to the RO.

Figure B.1 illustrates a number of points: the IP should always be between the RO and the ROP; however, the on-hand stock may fluctuate.

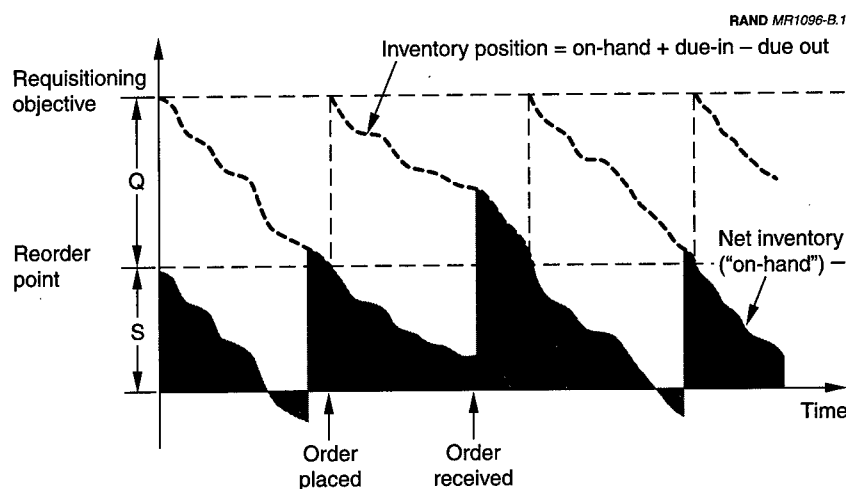


Figure B.1—An Illustration of the Relationship of the IP to the On-Hand Stock, the RO, and the ROP¹

tuates below the IP; and once the IP hits the ROP and a replenishment order is placed, the stock is at risk of stock-out and the risk is dependent on stocked depth at the ROP.² For example, in this figure the ROP was set for a high risk of stock-out (equivalently, a low service level) because, for the three complete *lead-time* cycles (replenishment order placement to order receipt) visible in the figure, the system went into back-order twice. If, however, as in Figure B.2, the ROP had been initially set high enough, the system might never have experienced a back-order status (S was increased by " d " amount to " $S+d$ "). Here we see sufficient stock exists to cover surges in demand after the on-hand quantity drops below the ROP.

¹Adapted from Hadley and Whitin (1963), Figure 4-1, p. 163.

²Actually, an item is *always* at risk of stock-out, even when the on-hand quantity is at the RO, given a large enough order. However, if the ROP is set properly, we tend to think of the stock being at risk of stock-out when the on-hand quantity is at or below the ROP.

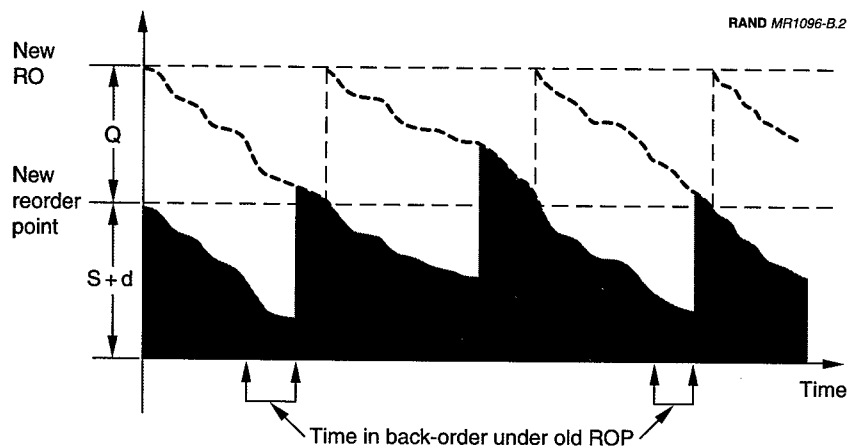


Figure B.2—An Illustration of the Effect of Increasing the ROP from “S” in Figure B.1 to “S+d”

The difference between the RO and the ROP is often referred to as the *operating level* (“Q” in Figures B.1 and B.2); for a given item, the larger the operating level the less often replenishment orders must be placed and the less exposure to risk of stock-outs. The ROP is said to comprise lead-time stock and safety stock (shown combined together as “S” in Figure B.1 and as “S+d” in Figure B.2); the *lead-time stock* is the quantity of a stock required to meet the average demand for the average lead time; the *safety stock* is additional stock kept for surges in demand and delays in lead time.

As shown in Figure B.3, the *lead time* is the time between reaching the ROP and when the subsequent replenishment order items are available to satisfy future demands. The OST is a subset of the lead time and is usually considered to be the time from when the order is sent to the wholesale system to when it is received locally. If replenishment orders are placed automatically when the ROP is reached, and if the time to restock the local inventory is negligible, then the lead time is equal to the OST.

The ROP should be set to provide an acceptably low risk of stock-out between the time when the stock level reaches the ROP and when the

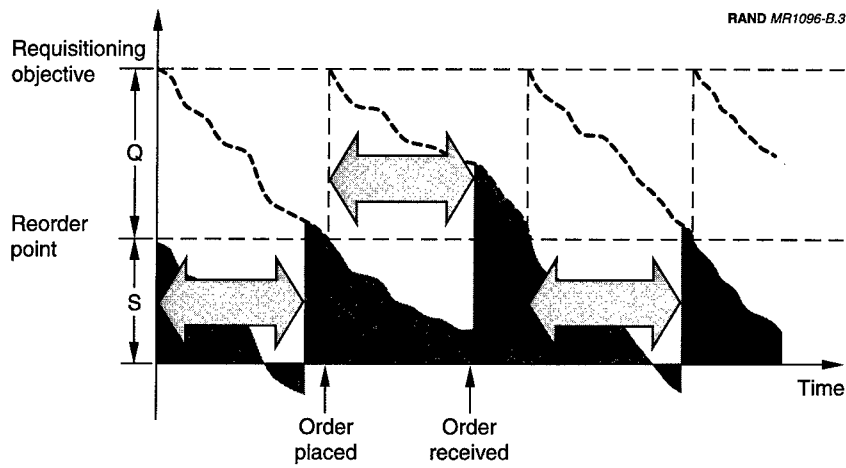


Figure B.3—Lead Time Indicated by the Gray Arrows.

order actually arrives. Setting the ROP is a trade-off between the cost of incurring one or more stock-outs and the cost of holding additional safety stock. This idea of “risk” is important and worth belaboring. Setting the ROP for an item is an exercise in gambling. The gamble is that if a replenishment order is placed as soon as the on-hand stocks reach the ROP that the replenishment order will come in before the on-hand stocks are used up during the lead time. That is the same as setting the ROP to be larger than the number of demands occurring during most of the lead times, assuming the inventory manager does not want to run out very often. The risk to the inventory manager is the chance that the safety stock will not last and the ROP should be set based on the maximum risk the inventory manager is willing to assume.

COMMON MARINE CORPS REORDER POINT AND THE OPERATING LEVEL MISCONCEPTIONS

Many misconceptions exist regarding how the current inventory system works, the purpose of specific components of the RO and ROP (such as the safety level), and the effects various changes have on the supply system. Some of these misconceptions are codified into Marine Corps policy documents, and others have come up during discussions. We will discuss a few of the more important ones here.

1. **Misconception:** The “operating level is that quantity of material required to sustain operations during the interval between the initiation of replenishment action and the arrival of successive shipments into the supply system.” (MCO, 1992, paragraph 1005, pp. 1-9, and Appendix A, p. A-6).

Fact: This definition applies to the ROP, which comprises the safety stock and the lead-time stock. The ROP is set so that material is available during the lead time required to replenish stocks.

2. **Misconception:** The operating level is set to provide the specified days of supply of materiel for deployment.

Fact: On-hand stocks can fluctuate anywhere between the RO and zero (and perhaps beyond if returns and back-orders are considered). The operating level, properly set, provides no assurance that items will be available at some point in the future. Only such stocks as the war reserves that are insulated from filling everyday demands can be guaranteed to be available.

3. **Misconception:** “The rationale for the 60-day OL is presumably to meet our doctrinal stand-alone sustainment capability.” (USMC Materiel Command Integrated Logistics Capability, 1999, p. A-21.)

Fact: As just discussed, the size of the operating level has nothing to do with sustainment. Generally speaking, the operating level only controls how often replenishment requisitions must be submitted.

4. **Misconception:** The operating level should be set to ensure that the RO is high enough that large batches and other causes of demand variability do not deplete the on-hand stocks to zero.

Fact: This notion confuses the operating level with safety stock and the RO with the ROP. The safety stock is designed to account for demand variability. Larger demand variability implies a larger safety stock and higher ROP; this only indirectly causes the RO to be larger.

5. **Misconception:** “. . . while certain low rate production or long-lead-time items may require higher operating levels.” (USMC Materiel Command Integrated Logistics Capability, 1999, p. A-21.)

Fact: This statement also confuses the operating level with safety stock and the RO with the ROP. Items with longer lead times require larger safety stock, which means a higher ROP *not* operating level.¹

6. **Misconception:** “The construction of stockage levels will be based on a combination of operating level (OL), actual order/ship time (OST) or procurement lead time (PLT), when available, and defined safety levels. This method of computing stockage levels should minimize total variable costs for any given supply performance or investment objectives.” (MCO, 1992, paragraph 1000.2, pp. 1–3.)

¹Technically speaking, an interaction between the operating level and the ROP affects the *overall* chance of stocking-out. But in Marine Corps’ systems where the operating level is relatively large, the primary buffer against stocking-out is the safety stock in the ROP.

Fact: The use of OL, SL, and OST may or may not minimize total costs depending on how they are set. Clearly, using the DOS methodology will not do this. Furthermore, using OST and not lead time will cause the Marine Corps to underestimate ROPs whenever automatic replenishment (that is, when replenishment orders are not issued when the ROP is reached) is not utilized as the norm.

7. **Misconception:** "Stock Objective. The maximum quantity of materiel authorized to be on hand to sustain current operations at any level of supply. It consists of the sum of stocks represented by the OL and the SL." (MCO, 1992, p. A-11.)

Fact: The stock objective, as "the maximum quantity of materiel authorized to be on hand," is synonymous with the RO (assuming no due-outs). As discussed in Misconception 6 above, and elsewhere in MCO P4400.151B, the stock objective is thus the sum of the OL, SL, and OST/PLT stock. Furthermore, it is incorrect to state that a stock objective (RO) can "sustain current operations at any level of supply." Calculations used to determine the RO must be based on past observed demand patterns and on an assessment of future demand, demand variability, and resupply availability; if these assessments are at variance with operational tempo and resupply then the RO may fail to support current operations.

USING THE BOOTSTRAP TO CALCULATE REORDER POINTS

Classic inventory theory models the distribution of an item's demands parametrically, making particular assumptions about the variability of demands and lead times when setting the ROP. Poisson, normal, and negative binomial distributions are commonly used. Such distributions allow simple calculation of the demand's variance so that the ROP can be specified in terms of a mean demand plus some multiple of the standard deviation. This approach works well if the distributional assumptions are appropriate. However, such distributional assumptions rarely apply to USMC data.¹

We address this problem with an empirically based resampling strategy, which uses bootstrapping techniques on historical demand data, to estimate the lead-time demand distribution. From this, quantiles for a given probability of stock-out can then be used to set ROPs. As discussed in any standard inventory theory text (see, for example, Tersine [1994] or Arrow, Karlin, and Scarf [1958]), there are three basic inventory problems, listed below in order of increasing complexity:

1. Fixed (deterministic) demand and lead time.
2. Stochastic demand, fixed lead time.
3. Stochastic lead time and demand.

¹See Fricker and Goodhart (forthcoming) for additional discussion and a comparison of the bootstrap with parametric techniques.

The first case will not be considered further as it is trivial to solve arithmetically (though we note here that the bootstrap methodology to be described applies equally well to this case as to the more complex, stochastic cases). For the second and third cases, let t denote time for some basic unit (days, weeks, etc.) and let D_t denote the number of items requested (the demand) during period t . Let L denote the lead time (in the same units as t) for an item's replenishment order and let r denote the probability that a stock-out occurs during the lead time. Assume that the demands D_t and the lead times L are both independent and identically distributed (*iid*) from some unspecified distributions, and that D_t and L are independent. Let

$$S(J,L) = \sum_{t=J+1}^{J+L} D_t$$

represent the cumulative demands occurring in L successive time periods after time J when the inventory position is first at or below the ROP. Denote $P\{S(J,L) \leq x\}$ as $F(x)$, the distribution function of lead-time demand (LTD). To achieve a particular probability of stock-out when the inventory is at the ROP, set $x = F^{-1}(1 - r)$ where x is the chosen ROP so that $P\{S(J,L) > x\} = r$. We refer to r as the *risk* or the probability of stock-out; the quantity $100(1 - r)$ is often referred to as the *service level*.

Both Bookbinder and Lordahl (1989) and Wang and Rao (1992) bootstrap observations from the LTD, assuming that actual observations for m lead times are available $\{S_1, \dots, S_m\}$. We do not use this approach for two reasons: our data do not have information about the inventory position, so we cannot identify the actual lead times, and, under the assumption that demands are independent of inventory position, only using data from the lead times ignores the majority of demand data. That is, if demands are independent of inventory position, then for any particular value of L it follows that $S(J,L)$ is equal in distribution to $S(j,L)$ for $j = 1, 2, \dots$,

and so we need not restrict the data to only the lead-times.

Under these assumptions, for a fixed lead time, the bootstrap (Efron and Tibshirani, 1993) can be used to create an empirical distribution

\hat{F} as follows. Suppose that for each NSN i a set of historical demands exists for N periods: D_1^i, \dots, D_N^i , where N is much larger than L . For a fixed lead time of L periods, the bootstrap methodology is applied by randomly drawing with replacement L observations from $\{D_1^i, \dots, D_N^i\}$ many, say M , times. The bootstrap observations are denoted $D_{(1)}^i, \dots, D_{(L)}^i$, with the iteration subscript suppressed for clarity. For each iteration, then, the bootstrap demand quantity for a lead time is calculated as

$$\hat{S}^i = \sum_{j=1}^L D_{(j)}^i.$$

The empirical distribution \hat{F}_i is calculated from M bootstrap statistics $\mathbf{S}^i = \{S_1^i, \dots, S_M^i\}$ and the ROP is then set choosing the quantile corresponding to a chosen risk r , $ROP_i = \hat{F}_i^{-1}(1 - r)$.

For the variable lead-time problem, the bootstrap methodology is further adapted by first randomly sampling from an empirical distribution of lead times and then resampling the number of demand periods corresponding to the chosen lead time. For example, given an empirical lead-time distribution for NSN i , a lead time is randomly selected; say l was chosen, then l periods are randomly drawn from the historical data, from which $\hat{S}^i(l)$ is calculated. As before, this process is repeated M times, where at each iteration a different lead time is randomly drawn from the empirical distribution.

Of course, since the empirical distribution \hat{F} will necessarily be discrete, some method will generally have to be employed to select the quantile which gives the risk closest (in some sense) to the specified risk. Efron and Tibshirani (1993) describe various methods for selecting quantiles which involve ranking the bootstrap observations and then choosing the $\alpha \cdot M$ observation (with rules for handling $\alpha \cdot M$ not integer). In our work here we set the ROP to the value in the empirical distribution with the risk closest to the desired risk in terms of Euclidean distance. That is

$$ROP_i = \{x: \min_{\forall x \in \mathbf{S}^i} |\bar{F}_i(x) - r|\},$$

where $\bar{F}_i(x) = 1 - \hat{F}_i(x)$. This is a more conservative approach than ranking, in the sense that ranking will always choose an ROP greater than or equal to our ROP, but our method generally requires less computer disk storage space because fewer than M observations can be stored when they are tied to bootstrap observations. Of course, other rules could be used to choose the ROP.

Some intuition is in order for why our modification to the Bookbinder and Lordahl (1989) and Wang and Rao (1992) methods works and is desirable. First, imagine a simple, ideal case in which the lead time for an item is deterministically one day and a large amount of historical demand data is available. Then, to evaluate the risk of stock-out for setting $ROP = x$ one would simply extract from the data all the demands for each day after the inventory position reached the ROP and count the fraction of times out of the total number of cases that the next day's demands exceeded x . This fraction would be a good estimate of the risk under the assumption that the distribution of future demands is similar to the distribution of past demands; the larger the amount of past data and the more similar the past and future demand distributions, the better the risk estimate. This is the motivation for constructing an empirical distribution of demands that occurred in the lead time after a replenishment buy order was placed, and it is the motivation for Bookbinder and Lordahl (1989) and Wang and Rao (1992) to bootstrap directly from the LTD.

However, under the assumption that inventory position and demands are independent we need not restrict the data to specific lead times, so we can generalize in this simple example by constructing the empirical distribution using all of the daily demands. Such an independence assumption is generally very reasonable to make for military supply systems, because demands are made without any knowledge of inventory position.

For lead times in excess of one day, we can further generalize by starting at the earliest day for which data are available and taking blocks of days sequentially that correspond to the lead time. This approach would allow us to use all of the demand data, rather than only the lead-time data, which are (typically) a small subset of the demand data. The difficulty with historical data, even when using all of the demand data, is that it is generally insufficient to construct a useful empirical distribution. For example, the Marine Corps yields

only two years of historical data; average lead times of 20 to 30 days make for too few LTD observations from which to construct a useful empirical distribution. Furthermore, even if a significant amount of historical demand data were available, the older data could be suspect in terms of distributional stationarity, so that it would be wise to limit the amount of older data utilized. Using the assumption of independent daily demands and independence between lead times and demands, we overcome this obstacle by resampling from the restricted historical data, bootstrapping "alternate" LTD observations that we use to capture the uncertainty related to observing limited demand data. From these alternative observations, we then create an empirical distribution, as in the previous simple example.

CALCULATING THE BOOTSTRAP ROPs FOR I MEF

We calculated the bootstrap ROPs for either a fixed or stochastic lead time using a resampling methodology as follows. One year's worth of I MEF requisition data (February 1, 1997, to January 31, 1998) was compiled into daily demands by NSN. That is, a matrix listing NSN by total number of items requisitioned daily was created from the I MEF General Account. We also created an empirical lead-time distribution for replenishment orders, for which, to simplify calculations and avoid problems with small sample sizes, we assumed that all NSNs had a common lead-time distribution. Then, for each NSN, the sum of a random number of days L (either fixed or stochastic) was selected (with replacement) 1,000 times, and an empirical distribution of the lead-time demand was created.

For the stochastic lead times, L was first randomly drawn from a lognormal distribution fit to actual lead-time data (truncated to a maximum of 30 days) taken from non-back-ordered replenishment OSTs.² Then, for a given risk r , the ROP was chosen as that value in the cumulative empirical lead-time demand distribution that was closest in terms of Euclidean distance to the specified risk. We performed the resampling using SAS. For approximately 35,000 NSNs with positive demands in the year period of interest, with 1,000

²Use of the lognormal distribution was a parametric simplification taken for computational convenience.

resamples drawn for each NSN, the program took approximately 45 minutes to run on an Ultra Sparc 5 with a 333-MHz processor.

For simplicity's sake the RO was defined, for those items with $ROP_i > 0$, as $RO_i = ROP_i + \max\{1, \text{round}(Z \text{DOS}_i)\}$, where DOS_i is the average daily demand for item i and the *round* function rounds the quantity to the next higher whole integer if the decimal fraction was greater than or equal to one-half, and it rounded down otherwise. (Thus, in the non-dollar banded scenarios, if $ROP_i = 0$ then $RO_i = 0$.) That is, the operating level was set to a fixed DOS, Z . (Other methods of setting the operating level, such as the economic order quantity, could be used, but will not be discussed here.) For the bootstrap GABFs the on-hand quantity was randomly set uniformly between the RO and the ROP, independently for each item. We refer to the bootstrap ROPs, ROs, and derived on-hand quantities as the *bootstrap GABF*.

Prior to evaluating the bootstrap ROPs for the full General Account using the Virtual SMU, as in Chapter Four, we evaluated the bootstrap ROPs' performance in a controlled comparison of desired and achieved service levels. Table D.1 shows the results of this comparison for five desired service levels and a fixed 30-day LTD. Here we took the entire I MEF GABF and calculated ROPs for each NSN, assuming a fixed 30-day lead time using data collected in the 12 months preceding February 1998, and evaluated them in five 30-day periods taken from calendar year 1998 (30 days was arbitrarily chosen simply to demonstrate the performance of the bootstrap ROPs by comparing them with the actual demands observed in new 30-day blocks of time³). The idea of the test is to assume that the stock levels for all stocked NSNs (defined as an NSN assigned a positive ROP by the bootstrap) have reached their ROPs simultaneously and then to evaluate how many of the NSNs stocked-out in a new 30-day period. For a given risk level and perfectly set ROPs, the expected number of NSNs experiencing a stock-out should be the risk level multiplied by the total number of NSNs.

³The "new" is important because we used the demand data from February 1, 1997, to January 31, 1998, to calculate the ROPs, so we did *not* use the same demand data to then test the methodology's performance. See Table D.1.

Table D.1
Comparison of the Probability of Not Stocking-Out ($1 - r$)
with the Actual Observed Rate for Five 30-Day Periods for All
Lines (NSNs) That Had a Positive Bootstrap ROP

Service Level ($1 - r$) \times 100%	Percentage of NSNs That Did Not Stock-Out for Five 30-Day Periods				
	March 1– March 30, 1998	March 31–April 29, 1998	April 30– May 29, 1998	May 30– June 28, 1998	June 29– July 28, 1998
80.0	85.7	69.6	85.5	88.1	84.5
85.0	89.2	77.5	89.0	91.4	88.8
90.0	92.4	85.0	92.6	94.1	92.7
95.0	95.4	91.0	95.3	96.4	95.6
99.0	97.6	96.0	97.7	98.3	98.2

Table D.1 shows that the 30-day fixed lead-time bootstrap ROPs track well for a given service level when compared with the actual demands observed for five 30-day periods. For all but one of the columns, the actual results are slightly conservative for the smaller service level values, most likely because of the decision rule for selecting the ROP from the empirical distribution. However, the second column (March 31–April 29, 1998) does show a period in which the demands were unusually high in comparison with the past year's demands (upon which the bootstrap calculations were based). This result demonstrates that any methodology, no matter how good, that uses the past to plan for the future cannot perfectly predict that future.

Also note in Table D.1 that, as the risk becomes very small, roughly $r < 0.05$, the percentage of NSNs not stocking-out falls below the specified service level. This is because at very small risk levels the bootstrap essentially reduces to the sum of the maximum 30 days of demands observed in the past year, and, again because the past is not always a good indicator for the future, some NSNs experience demands in the new period greater than the previous year's maximum. However, as the table shows, at small risk levels the difference between desired and actual performance is quite small—usually within a percentage point or two.

Overall, Table D.1 demonstrates that the bootstrap methodology using independent sampling by days works well, from a practical standpoint, for this data. Other calculations (not shown here) that increased the number of resamples per NSN from 1,000 to 10,000 and sampled by weeks instead of days did not yield any appreciable difference in results.

GENERALIZATIONS TO THE BOOTSTRAP ROP

Our application of the bootstrap might seem to differ from its more common use in estimating the standard error of a mean. To illustrate the more usual application, consider a set of observations Y_1, Y_2, \dots, Y_n and their mean,

$$\bar{Y} = \frac{1}{n} \sum_{k=1}^n Y_k.$$

To estimate the standard error of \bar{Y} , one applies the bootstrap by resampling with replacement M sets of n samples from Y_1, Y_2, \dots, Y_n : $Y_{(1,k)}, Y_{(2,k)}, \dots, Y_{(n,k)}$, $k = 1, \dots, M$. From these resamples,

$$\bar{Y}_{(k)} = \frac{1}{n} \sum_{j=1}^n Y_{(j,k)},$$

is calculated for each k , and then $s.e.(\bar{Y})$ is estimated as

$$\widehat{s.e.(\bar{Y})} = \left[\frac{1}{M-1} \sum_{k=1}^M (\bar{Y}_{(k)} - \bar{Y}_{(.)})^2 \right]^{\frac{1}{2}},$$

where

$$\bar{Y}_{(.)} = \frac{1}{M} \sum_{k=1}^M \bar{Y}_{(k)}.$$

Depending on the statistic being estimated, $\bar{Y}_{(i)}$ may be used as the point estimate, or some other, more standard statistic. In this example, \bar{Y} would probably be used as the point estimate of the mean, and the bootstrap standard error,

$$\widehat{s.e.(\bar{Y})},$$

as the measure of the point estimate's variability.

The key ideas in the bootstrap are (1) resampling n observations with replacement from the original n , and (2) thinking of these sets of bootstrap observations as "new" data that are representative of the underlying probability distribution that generated the original data. In our application, a computational methodology equivalent to the one we described in the first part of the appendix is as follows:

1. Resample with replacement N observations from $\{D_1^i, \dots, D_N^i\}$ to generate the bootstrap observations $D_{(1)}^i, \dots, D_{(N)}^i$.
2. Then generate a lead-time observation l (stochastically or deterministically) and choose l days randomly, with replacement, from $D_{(1)}^i, \dots, D_{(N)}^i$.
3. Repeat step 2 M times, calculate \hat{F}_i from $\{\hat{S}_1^i, \dots, \hat{S}_M^i\}$, and then choose the quantile corresponding to the desired risk r , $ROP_i = \hat{F}_i^{-1}(1 - r)$.

This represents one bootstrap ROP for item i , which we could use as the point estimate of the actual distribution quantile. To calculate the standard error for this point estimate, we would repeat steps 1–3 many times and use those replications to compute the standard error in the usual way described above.

To see the equivalence between steps 1–3 and the previous Monte Carlo-like description, note that drawing a random sample (with replacement) of lead-time demands in step 2 from the bootstrap distribution, which itself was drawn randomly with replacement in step 1, is exactly the same as drawing a random lead-time number of demands from the original set of demands. That is, the first two steps simply collapse into one, in which l demands are randomly drawn with replacement from the original set of N daily demands.

However, under autocorrelation or some other type of data structure assumption, the bootstrap demands resample in step 1 might be conducted differently. For example, we might bootstrap weeks or months of demands and randomly assemble the weeks/months into a new demand stream. Similarly, with autocorrelation we also might choose contiguous blocks of demands to calculate the LTD in step 2. Other techniques can also be employed to account for data structure; see Efron and Tibshirani (1993, Chapter 8) for additional discussion.

Specifically, under these more complicated scenarios, steps 1 and 2 do not reduce to the simple Monte Carlo approach previously described. Further, if the empirical lead-time distribution consists of limited data, one might also add a step between 1 and 2 to bootstrap the lead-time distribution. Thus, it is useful to think about these types of problems in the more general bootstrap framework. This framework helps clarify the problem by splitting it into two sub-problems: incorporating the uncertainty stemming from limited observed data (either demand, or lead time, or both) using bootstrap resampling, and estimating the lead-time demand distribution via Monte Carlo or some other appropriate means.

For seasonal items, the key modification is to restrict the demand stream being resampled to the particular season of interest; if it is a summer item, only use the summer demand stream to estimate the LTD distribution. For other situations, the resampling strategy should be appropriately modified. For example, in the Marine Corps, the bootstrap methodology can be generalized to more complicated inventory problems, such as stocking repairable items at the RIP, where an issued item is coupled with a returned broken item that may be repaired or refurbished and returned to stock.

In spite of this more complicated scenario, a bootstrap scheme can be readily implemented that would account for these additional requirements and still produce a simple risk distribution from which to set ROPs. To do so, two resamples must be drawn for each sampled lead time, one representing the demands and one representing repaired items returned to stock. The net positive number of demands after subtracting the repaired items could then be used to estimate the LTD for new items. A similar approach could be used for the Marine Corps ICP as well.

Finally, estimation of the LTD distribution could be made more sophisticated through the application of kernel density estimation techniques. That is, rather than use the simple empirical distribution generated by the bootstrap, a discrete kernel function could be applied to smooth the density estimate, should such smoothing be found advantageous. As these ideas illustrate, there are many ways to apply the bootstrap methodology to estimate lead-time distributions. Each variation must be judged on its applicability to the particular organization and problem.

Appendix E

THE VIRTUAL SMU

Chapter Three presented a number of concepts for improving Marine Corps inventory performance and demonstrated their benefits independently and sometimes abstractly. Yet, in practice these techniques will be used in concert and with other policies and procedures already in effect. It is thus very desirable to analyze the performance of these techniques in as real a setting as possible—a setting that takes into account how demands actually arrive at the SMU and other “real world” issues.

In order to achieve this, RAND developed the “Virtual SMU,” which, at its most fundamental level is simply an inventory accounting tool that replays a past sequence of actual demands and calculates all of the standard inventory metrics, including ERO fill rate. The idea of such a tool is that various policy alternatives can be evaluated under the rigors of actual demand streams with the goal of selecting inventory management policies, techniques, and algorithms that perform well. It seeks to answer the question, “How can the SMU achieve the best possible *overall* performance (in terms of performance metrics such as fill rate and ERO fill rate) within imposed funding constraints?”

Analyzing the effect of inventory policy changes prior to their implementation is a desirable goal, though one not routinely achieved in practice for the Marine Corps. For a complex metric such as the ERO fill rate, this desire becomes necessity, however, because it is not obvious how various policy changes will affect the ERO fill rate. Furthermore, it is not clear a priori if and how changes

designed to improve the ERO fill rate will affect other aspects of SMU operations and other inventory metrics.

WHAT IS THE VIRTUAL SMU?

The Virtual SMU is open software code, currently written in SAS, which performs all of the standard inventory accounting procedures that an inventory management system would. It starts with an initial inventory balance, subtracts demands and/or accounts for due-outs as requisitions are presented, orders and adds replenishments as needed, and otherwise tracks financial constraints and calculates inventory performance metrics. While simple in concept, the code is fairly complicated because it allows for a variety of policy alternatives. Furthermore, it is "open" in the sense that it is uncompiled to allow further policy options to be entered as modifications to the code.

Basic Operation of the Virtual SMU

The general idea is to start out at some IP, play demands against the inventory for a period of time, and then calculate the performance of the system. The overarching goal is to use as little simulated data as possible, only resorting to simulation if actual data are not available. A description of the basic steps of the Virtual SMU is as follows:

- The Virtual SMU starts with ROs and ROPs from either an existing GABF or an alternative GABF, such as one created using the bootstrap and dollar-banding techniques. The on-hand stock can be set either to the actual on-hand stock from the existing GABF or through some other mechanism, such as setting on-hand stock randomly between the RO and ROP for each NSN. Similarly, the due-in and due-out can be set to the actual GABF values or zero. Setting the due-in and due-out to zero while setting the on-hand stock randomly between the RO and ROP can be thought of as starting the system with a "clean slate." Using the actual values allows an analysis to consider the transition problem: how to get from the existing inventory to a new IP.
- Next, a particular period's worth of "new" demands is specified, say for three months from May 1 through July 31, 1998. These

demands are then “played back” day-by-day through the Virtual SMU, subtracting each demand from the on-hand stock if available, adding to the due-out if not, or passed through to the wholesale system in accordance with standard Marine Corps practices.

- Each day, the IP is computed and tracked for all NSNs. Currently, replenishment buys are placed automatically, weekly, by selecting NSNs with IP less than or equal to the ROP, ranking them according to some prespecified criteria, and placing replenishment buys subject to some specified fiscal constraint. Due-in dates are assigned (stochastically or deterministically) for the replenishment requisitions, and they are added to the on-hand stock as they arrive.
- Inventory statistics are maintained over the course of the simulation, and overall performance metrics are calculated at the conclusion of the simulation period.

Details of the Virtual SMU

Figure E.1 gives a schematic outline of the various parts of the Virtual SMU. We will discuss each in some detail below.

Data Inputs. Data inputs to the Virtual SMU come from two primary sources, actual demand data and IP information from an actual GABF. The demand data can come from any source that captures General Account requisitions in terms of NSN, quantity ordered, and date submitted. The demand stream should consist only of consumable items that would be requisitioned from the General Account and must not include replenishment requisitions from the SMU.^{1, 2}

¹Assuming the analysis is on the General Account. If the analysis is related to some other inventory, then, of course, the demand stream for that inventory must be used.

²For the simulations in Chapter Four, we began with all requisitions that had an AE1 and D6T transaction in SASSY. From these we deleted all SAC 2 and 3 requisitions, all replenishment requisitions (defined as requisitions having a UAC of MMCXXX), and all float items (defined as requisitions having a UAC of MMFXXX). We also removed food items (defined as requisitions having an FSC of 8970), because these items are funded separately.

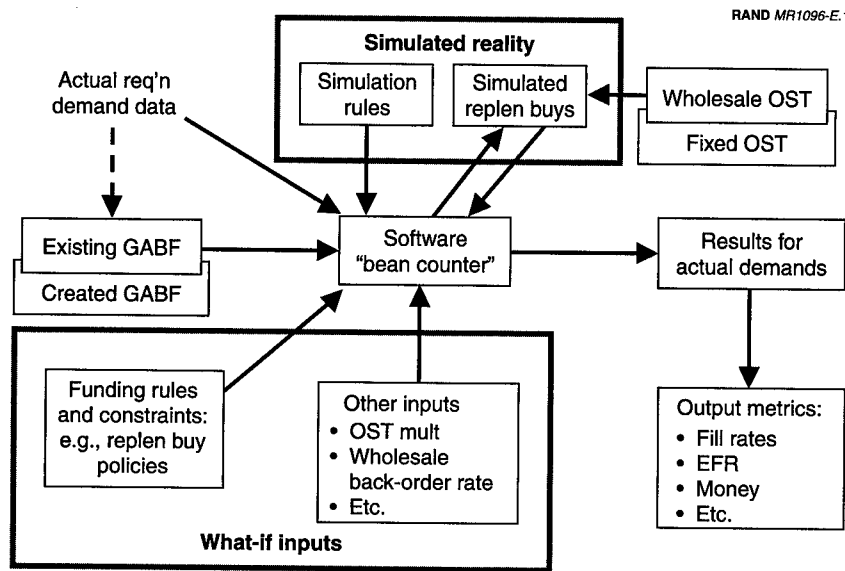


Figure E.1—Schematic Outline of the Virtual SMU

GABF inputs consist of the RO, ROP, unit price, fixed-level, and on-hand quantities by NSN for the entire inventory. These quantities may come from an actual GABF or may be created with some other mechanism, such as the bootstrap ROP. They may also be combined to explore how to make the transition from an existing IP to a new inventory by taking the on-hand from an existing GABF and newly created/defined ROs and ROPs. In the example to follow at the end of this appendix, we started from an existing IP from an actual GABF. The work in Chapter Four combines existing and new data. The fixed-level quantity (essentially materiel set aside as war reserve) was set to zero for all NSNs in the work to date, but this could be set to actual or desired levels to explore the effects of war reserve on inventory performance.

A third possible input of actual data comes in the form of "actual" wholesale (non-back-ordered) lead times. The word "actual" is in quotation marks because the current coding uses a stochastic draw from an empirical distribution built from actual wholesale OSTs.

This feature exists to explore the effects of randomness in wholesale lead times on system performance. The empirical distribution can be changed to explore the effects of improving or degrading lead times, and the lead time can also be set to a fixed, deterministic value.

“What-If” Inputs. What-if inputs are those things related to policy changes that the SMU, the Marine Corps, and DLA can change to affect inventory. For the SMU and Marine Corps this includes the GABF (RO and ROP) and such policy alternatives as replenishment buying policies (frequency and apportionment of funding); for DLA it includes such adjustments as the wholesale back-order rate.

Creating a new GABF has been previously discussed. A number of replenishment buy policies were coded to reflect possible strategies for automatically executing buys. For the Virtual SMU, it was impractical to mimic the time-consuming actual practice of manually scrubbing recommended replenishment buys. Indeed, one reason for creating the Virtual SMU was to evaluate automatic buy policies to eliminate the need for significant manual effort in the replenishment buy process.

Currently the automatic buy algorithm has been coded as follows. Schedule buys at predetermined times, say once a week. At that time, rank all NSNs with IPs less than or equal to the ROP by unit price.³ Start with the least expensive item and buy the quantity to return the IP to the RO. Continue down the ranked list until the week’s funding constraint is reached. There are two funding alternatives for the buy policy. In the “average” policy, the funding constraint is computed by taking the quarter’s General Account funds, setting aside a fixed percentage to cover pass-throughs,⁴ and then dividing the remainder by the number of weeks in the quarter. Unspent replenishment buy funds are carried over into successive periods. In the “conservative” policy, the unspent funds are not car-

³Alternative ranking mechanisms may be used, such as CEC or average requisition priority.

⁴Requisitions for items not stocked by the SMU are “passed through” to the wholesale system. Similarly, high-priority requisitions for stocked items that the SMU has run out of are passed through. The SMU must reserve sufficient funds to pay for these pass-throughs as they occur.

ried over into subsequent replenishment buys but are instead put into the pass-through funds.

Another what-if input is the wholesale back-order rate. This rate defines the fraction of the NSNs that will be in wholesale back-order status at some time during the simulation. Currently, it randomly selects 8 percent of the NSNs with a random back-order start date. From this random date, the NSN is defined to be in back-order status for 150 days. If an order is received from the SMU during this time then the replenishment order is not shipped until the item comes out of back-order status. This results in approximately 5 percent of the NSNs being in wholesale back-order on any given date in the simulation—approximately what occurs in practice.

Simulated Reality. Simulated reality is so titled because it is here that simulation must be used to account for unobserved reality. For example, when an automatic replenishment buy is placed we must simulate how long that replenishment buy takes to return to the SMU. Some of that is based on actual wholesale OST, and some is based on the simulated back-order rate previously discussed. Also, rules must be programmed into the “bean counter” portion of the software to account for unavailable information, such as the order in which a day’s (actual) requisitions were filled. In the case of the Virtual SMU, high-priority requisitions are always filled first, followed by low-priority items. Within each priority, the requisitions are filled in NSN order. Neither of these rules is likely to have always (or even sometimes) been followed in practice, but no other information is available.

The “Bean Counter” and Results for Actual Demands. The “bean counter” is simply an accounting device that takes requisitions that come in from the actual demand data stream, subtracts each demand from the on-hand stock if available, adds to the due-out if not, or passes it through to the wholesale system in accordance with standard Marine Corps practices. At the start of each day it first checks to see if any replenishment requisitions “arrived” and, if so, first adds those quantities to the appropriate on-hand quantities. Once a week, it executes an automatic replenishment buy. The results are tracked in the form of how the requisitions were filled—out of the General Account stocks or from wholesale.

Output Metrics. Multiple metrics are output, including accommodation, satisfaction, and fill rates, as well as the ERO fill rate. Weekly fill rates by high- and low-priority requisition are output. The starting and ending IP is characterized in terms of number of lines stocked and the percentage of lines with the IP less than the ROP or greater than the RO. Financial metrics are also reported, including the percentage of funds spent on replenishment requisitions and pass-throughs.

Assumptions and Limitations in the Current Version

The Virtual SMU currently uses a number of simplifying assumptions. These were applied under the modeling paradigm of avoiding unnecessary complexity. That is, we avoided introducing details thought to be of a secondary nature to the fundamental questions related to the bootstrap ROP under study. However, once the basic issues related to properly setting the ROs and ROPs are addressed, future work may want to look into those secondary effects and thus include the following in the Virtual SMU:

- The Virtual SMU does not currently account for returns to the SMU. Returns certainly affect IP and are a primary cause for the NSNs with IP greater than the RO that have been observed.
- Related to returns, the Virtual SMU currently does not distinguish the effects of MEU⁵ operations on the SMU, such as parts block creation and returns after deployment. It is likely that these operations affect inventory levels and should be separated so that the effect of changing policies for the MEUs can be evaluated.
- The Virtual SMU does not account for initial inventory provisioning (IIP) for newly fielded systems. At the moment, it treats all NSNs alike, as previously described.

⁵MEU is an acronym for Marine Expeditionary Unit which, as the name suggests, is a mobile unit usually embarked on a ship. MEUs deploy with a set of repair parts known as a "parts block." The parts block is assembled prior to deployment and is designed to support the type of mission or missions that the MEU is deployed for or may expect. Because these repair parts are assembled to reflect expected repair needs under the mission of the MEU, they will not necessarily be used during the deployment and often do not reflect the usual repair parts usage patterns.

- Similarly, the Virtual SMU currently does not account for interchangeable or substitutable NSNs. Rather, each NSN is assumed to be separate and unique.
- The Virtual SMU also does not allow for hot-item back-order.⁶ Rather, it fills all requisitions in the order they were requested.
- The Virtual SMU has been coded to allow for war reserves in terms of the GABF fixed-level. However, in the simulations so far the fixed-level has been set to zero. Additional work to study how and at what level to set the fixed-level, and good policies for managing these assets need to be evaluated.
- ERO fill rate is currently calculated only for the input demand stream. Thus, for example, the simulations in this report do not include repairable items from the RIP.

Practical Matters

The Virtual SMU was written as an analyst's tool and, as such, no effort has been made to make it "user friendly." Rather, users of the program must have a solid understanding of SAS and must be sophisticated enough to understand the potential pitfalls and limitations of such software. The Virtual SMU was specifically written in this format, however, to facilitate the analysis of new policy alternatives. That is, because it is impossible to foresee and code all possible alternative inventory policies, the code must be left open for modification.

The program takes about 45 minutes to run on an Ultra Sparc 5 workstation for eight months of data for an MEF. This run-time is exclusive of any time it takes to assemble the input data, such as the demands or the GABF.

⁶A hot item back-order occurs when the SMU fills the demand for a "hot item" with the next receipt of that type of item from the wholesale system, overriding the original destination of the order.

IDENTIFYING EFFECTIVE SIMPLE POLICY ALTERNATIVES

The use of the Virtual SMU is easily illustrated with a simple example. Inventory managers at the SMUs have the daunting task of periodically evaluating replenishment buy recommendations for financial feasibility. Weekly, or perhaps more often, the inventory manager is confronted with a list of NSNs with IPs less than their ROPs, and the manager must decide which items and how much of each item to buy within that week's fiscal constraints. Of course, in the ideal inventory theory world, all of the items would be purchased to return the IPs back to the ROs. However, for a variety of reasons, in the real world of the SMU this is not always possible simply because the SMU cannot afford to buy the specified amount of all the items.

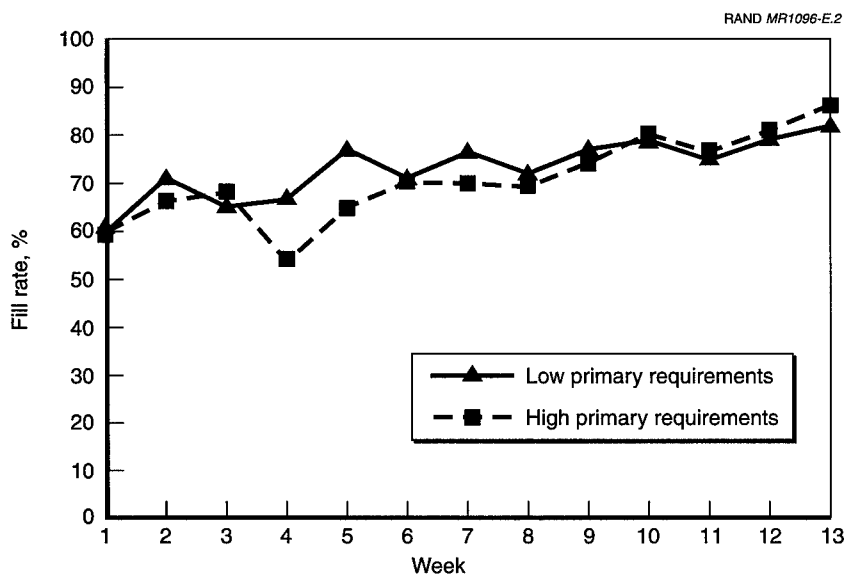
As Chapter Three discussed, much of this difficulty may result from inaccurate ROPs and ROs, and with the implementation of the bootstrap ROPs some of this trouble may be relieved. However, with such improvements yet to be implemented, the inventory manager is interested in what can be done today to alleviate the extensive manual effort that goes in to scrubbing the replenishment buys.

Using the Virtual SMU, we discovered that a very simple replenishment buy rule made significant improvements in inventory performance. That is, not only does the automatic replenishment buy rule relieve the inventory manager of the burden of a manual process, it seems to generate an inventory that performs better than the previous subjective method. The simple buy rule is to rank weekly all NSNs with IPs below their respective ROPs, from least expensive to most expensive. Then, starting with the least expensive, buy the quantity necessary to return the IP to the RO, proceeding through the ranked NSNs until the week's allocated funds are expended.

To demonstrate this, we employed the Virtual SMU as follows. We took the actual I MEF GABF for April 27, 1998, which included the actual ROs, ROPs, and on-hand stock amounts. For demands to the SMU we "played back" the actual requisitions presented to the SMU for the period May 1 to July 31, 1998. We gave the SMU a \$7 million budget for the quarter, which is roughly the actual quarterly budget, and set aside 60 percent of that to cover pass-throughs. We specified that replenishment buys would be placed every seven days and allocated one-thirteenth of the available funds to be spent each week.

That is, each week the Virtual SMU could spend slightly more than \$215,000 for that week's replenishment buy.⁷

We then ran the Virtual SMU for the quarter, allowing it to make automatic replenishment buys for the quarter—no human intervention occurred. Figure E.2 demonstrates the results. The figure shows an excellent upward trend in fill rate over the course of the quarter, starting in the low 60 percent range (which was close to actual) and progressing to above 80 percent by the end. The graph shows an initial three- to four-week lull (weeks one to four) because a fixed 18-day OST was programmed into this simulation and none of the due-ins from the actual GABF was carried forward, so for the first three weeks *no* replenishment orders were received.



**Figure E.2—Dramatically Improved Fill Rate Performance
Under a Simple Buy Policy**

⁷Actually, the algorithm was slightly more complicated than this because unspent replenishment buy funds could be carried over to the next period. So the first buy started out with a ceiling of \$215,000 to spend. If it all was spent then the next buy would have the same ceiling. If it was not all spent then the next buy would have one-thirteenth of 40 percent of the remaining funds.

Figure E.3 shows that the inventory rapidly adjusted to bring itself into proper alignment over the course of the quarter, too. In particular, the middle line shows that the inventory started out significantly out of position with respect to the number of lines with the IP less than or equal to the ROP—almost 7,000 lines. However, by the end of the quarter this was reduced to only 3,000, with evidence that this downward trend would have continued had the simulation continued past the quarter. Similarly, the number of lines with IP greater than the RO and the number of lines with a zero RO but stock on-hand ($IP > 0$) showed gradual decreases as stock was slowly used. Over the course of the quarter, the average fill rate was 72 percent and the ERO fill rate was 54 percent, both of which were up slightly from the actual performance of the MEF. However, these averages are deceiving because the performance at the end of the quarter was significantly better. Continuation of this buy policy into the next quarter could reasonably be expected to show fill rates of more than 80 percent and ERO fill rates in the 60 percent range. And these gains were achieved as a result of a simple automated buy policy *under the*

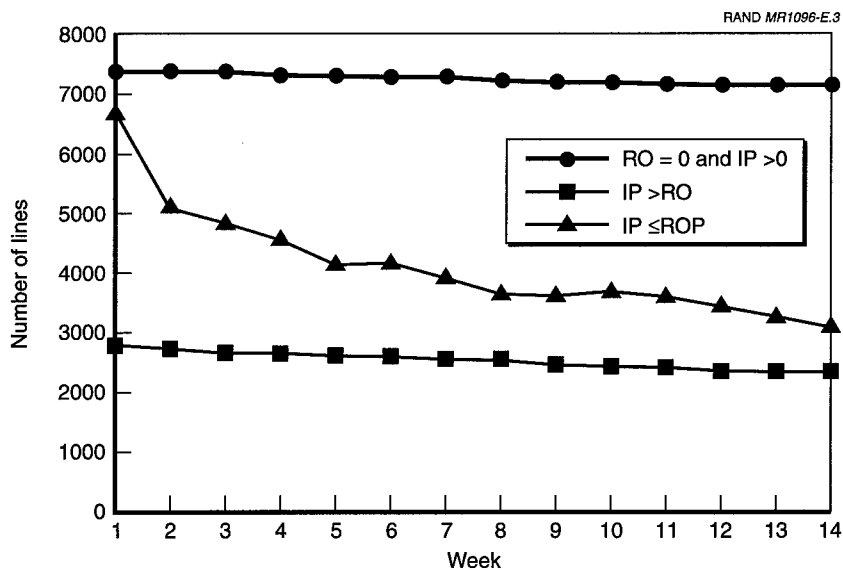


Figure E.3—Performance of the Inventory Under the Simple Buy Policy

existing, imperfect ROs and ROPs. Thus, not only does this buy policy relieve the drudgery and wasted man-hours of scrubbing replenishment buy lists, it also improves inventory performance.

APPLICATIONS OF THE VIRTUAL SMU

While the Virtual SMU was conceived of as a way to evaluate the performance of other, more complicated changes in inventory management practice, such as the bootstrap ROP and dollar-banding methodologies, the simple buy policy example demonstrates the power of a tool like the Virtual SMU.

The simple buy policy was originally conceived of as a “place holder” in the code until a more sophisticated policy was devised. No one expected it to perform well and, while a better policy may yet be identified, this simple buy policy is a significant improvement over current practice. It is a policy that is simple to implement, is simple to understand, and provides significant inventory performance improvement within current practice. Yet, without the evidence provided by the Virtual SMU output, its benefits would never have been convincingly demonstrated.

The Virtual SMU provides the Marine Corps with a sophisticated “what-if” inventory analysis capability that currently does not exist. It is a means of evaluating inventory procedures off line without resorting to experimenting on the field. As such, and as Chapter Four demonstrates, the Virtual SMU provides the Marine Corps an analytical means to test new supply policy and quantify the effects prior to implementation.

REFERENCES

- Arrow, K. J., S. Karlin, and H. Scarf, *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford University Press, 1958.
- Bookbinder, J. H., and A. E. Lordahl, "Estimation of Inventory Re-Order Levels Using the Bootstrap Statistical Procedure," *IIE Transactions*, Vol. 21, No. 4, 1989, pp. 302-312.
- Dumond, J., R. Eden, and J. Folkeson, "Velocity Management: An Approach for Improving the Responsiveness and Efficiency of Army Logistics Processes," Santa Monica, Calif.: RAND, DB-126-1-A, 1995.
- Edwards, T. J., and R. Eden, "Velocity Management and the Revolution in Military Logistics," *Army Logistician*, Vol. 31, No. 1, January-February 1998.
- Efron, B., "Bootstrap Methods: Another Look at the Jackknife," *Annals of Statistics*, Vol. 7, 1979, pp. 1-26.
- Efron, B., and R. J. Tibshirani, *An Introduction to the Bootstrap*, New York: Chapman & Hall, 1993.
- Fricker, R. D., Jr., and C. A. Goodhart, "Using a Bootstrap Approach for Setting Reorder Points in Inventory Systems," *Naval Research Logistics*, forthcoming.
- Hadley, G., and T. M. Whitin, *Analysis of Inventory Systems*, Englewood Cliffs, N.J., Prentice-Hall, 1963.

Krulack, Gen. Charles C., white letter on Precision Logistics, January 16, 1997.

Marine Corps Order P4400.151B, *Intermediate-Level Supply Management Policy Manual*, July 9, 1992.

Robbins, M. L., P. M. Boren, R. Eden, and D. A. Relles, "Measurement of USMC Logistics Processes: Creating a Baseline to Support Precision Logistics Implementation," Santa Monica, Calif.: RAND, DB-235-USMC, 1998.

Silver, Edward A., *Decision Systems for Inventory Management and Production Planning*, New York: John Wiley & Sons, 2nd Edition, 1985.

"Speeding the Flow: How the Army Cut Order-and-Ship Time," Santa Monica, Calif.: RAND, RB-3006, 1998.

Tersine, Richard J., *Principles of Inventory and Materials Management*, PTR Prentice Hall, 4th Edition, Englewood Cliffs, N.J., 1994.

U.S. Marine Corps Materiel Command Integrated Logistics Capability, "Case Study & Appendices A-C," March 4, 1999.

Wang, M., and S. S. Rao, "Estimating Re-Order Points and Other Management Science Applications by Bootstrap Procedure," *European Journal of Operational Research*, Vol. 56, 1992, pp. 332-342.